

Parallel Cleaning of a Network with Brushes *

Serge Gaspers[†] Margaret-Ellen Messinger[‡] Richard J. Nowakowski[‡]

Paweł Prałat[‡]

[†]Department of Informatics, University of Bergen,
N-5020 Bergen, Norway

[‡]Department of Mathematics and Statistics, Dalhousie University,
Halifax, NS, B3H 3J5, Canada

Abstract

We consider the process of cleaning a network where at each time step, all vertices that have at least as many brushes as incident, contaminated edges, send brushes down these edges and remove them from the network. An added condition is that, because of the contamination model used, the final configuration must be the initial configuration of another cleaning of the network. We find the minimum number of brushes required for trees, cycles, complete bipartite networks; and for all networks when all edges must be cleaned on each step. Finally, we give bounds on the number of brushes required for complete networks.

1 Introduction

Imagine a network of pipes with a biological contaminant that is: (1) relatively mobile so that contamination can spread from one area to another; and (2) that regenerates. For example, water pipes that have algae or zebra mussels contamination [5, 6]. To prevent the network from clogging, the pipes must be cleaned on a regular basis and steps must be taken to prevent recently cleaned pipes from being recontaminated. Because of the regeneration of the contaminant, the final configuration is the basis for the start of the next round of network cleaning. There is much work on models of contamination of networks where condition (1) is paramount, see [1, 8] for examples. Condition (2), regeneration, is new in this context. See [12] for a situation where the contaminant regenerates but is immobile.

The *(sequential) brush cleaning* model was introduced in [7, 10]. Initially, every edge and every vertex of a network is dirty and a fixed number of brushes start on a set of vertices. A vertex may be ‘cleaned’ if it contains as many brushes as dirty incident edges. At each step, one vertex v is cleaned and each incident, dirty edge is traversed by a brush from v . Once a brush has traversed a dirty edge, that edge has been cleaned and for this round of cleaning, may be regarded as being deleted from the network. Thus, a network has been cleaned once every edge has been cleaned. The brush number of network G , is the minimum number of brushes needed to clean G and is denoted $b(G)$. The model requires that the cleaning process

*Partially supported by grants from the ACEnet, NFR, NSERC, MITACS, and SHARCNET.
serge@ii.uib.no, messnger@mathstat.dal.ca, rjn@mathstat.dal.ca, pralat@mathstat.dal.ca

be *continual*: once the network has been cleaned, the network is regarded as contaminated again and the brushes re-clean the network. If the vertices are cleaned sequentially, it was shown in [10] that the process is always continual, or more specifically ‘reversible’. That is, let ω_0 be an initial configuration of brushes that will clean a network G , which leaves a final configuration ω_n of brushes. Using the initial configuration $\tau_0 = \omega_n$ of brushes, G can be cleaned leaving the final configuration $\tau_n = \omega_0$. The sequential brush number has also been studied on random regular networks [2, 11] and on random networks [13], showing a that the brush number is almost surely close to $dn/4$ for networks of (large average) degree d .

In this paper, we consider a variant of the cleaning model, introduced in [7] where at each step, every vertex which may be cleaned, is cleaned simultaneously. In this variant, the **parallel cleaning model**, the minimum number of brushes needed to clean a network G , with parallel firing of vertices, is called the *parallel brush number* and is denoted $pb(G)$. See Figure 1 for an example. It is also clear that the parallel cleaning process for this network is always reversible. This is not the case for the network G' in Figure 2, with one brush




Figure 1: The parallel cleaning process for a network G with one brush initially at each of v_1, v_6 .

initially at each of v_1, v_7 . It can only be cleaned once. The final configuration of brushes (one at each of v_4, v_5) is not a viable initial configuration of brushes if the edges of the network were to become dirty again. This provides motivation for the results of this paper. We wish to determine the minimum number of brushes needed to ensure a network can be parallel cleaned *continually*. Note that, although the practical situation being modeled only requires




Figure 2: The parallel cleaning process for a network G' with one brush initially at each of v_1, v_7 .

every edge to be cleaned for the network to be deemed clean, the assumption taken in [10] and taken here, is that a network has been cleaned once every vertex (and hence every edge)

has been cleaned. Although this viewpoint may seem unnatural, it simplified much of the analysis in [10]. Therefore, in Figure 1, at time $t = 3$, vertex v_4 is cleaned (although it has no dirty incident edges).

Section 2 formally defines the parallel cleaning process and the notion of continually parallel cleaning a network. In Section 3, we determine the exact number of brushes required to continually one-step clean a network (in each cleaning, all edges of the network are cleaned in the first step). Section 4 focuses on the minimum number of brushes to continually parallel clean various networks, finding exact results for cycles, trees and complete bipartite networks. Finally, Section 4.4 bounds the number of brushes needed to continually parallel clean a complete network K_n between $\frac{5}{16}n^2 + O(n)$ and $\frac{4}{9}n^2 + O(n)$, a table of exact values is also given.

2 Definitions

Following the terminology used for the (sequential) cleaning model in [10], at each time step t , $\omega_t(v)$ denotes the number of brushes at vertex v and D_t denotes the set of dirty vertices. An edge $uv \in E$ is dirty if and only if both u and v are dirty: $\{u, v\} \subseteq D_t$. Finally, let $D_t(v)$ denote the number of dirty edges incident to v at step t :

$$D_t(v) = \begin{cases} |N(v) \cap D_t| & \text{if } v \in D_t \\ 0 & \text{otherwise.} \end{cases}$$

Definition 2.1 *The **parallel cleaning process** $\mathfrak{C} = \{(\omega_t, D_t)\}_{t=0}^K$ of an undirected network $G = (V, E)$ with an **initial configuration of brushes** ω_0 is as follows:*

- (0) *Initially, all vertices are dirty: $D_0 = V$; set $t := 0$*
- (1) *Let $\rho_{t+1} \subseteq D_t$ be the set of vertices such that $\omega_t(v) \geq D_t(v)$ for $v \in \rho_{t+1}$. If $\rho_{t+1} = \emptyset$, then stop the process ($K = t$), return the **parallel cleaning sequence** $\rho = (\rho_1, \rho_2, \dots, \rho_K)$, the **final set of dirty vertices** D_K , and the **final configuration of brushes** ω_K*
- (2) *Clean each vertex $v \in \rho_{t+1}$ and all incident, dirty edges by traversing a brush from v to each dirty neighbour. More precisely, $D_{t+1} = D_t \setminus \rho_{t+1}$; for every $v \in \rho_{t+1}$, $\omega_{t+1}(v) = \omega_t(v) - D_t(v) + |N(v) \cap \rho_{t+1}|$; for every $u \in D_{t+1}$, $\omega_{t+1}(u) = \omega_t(u) + |N(u) \cap \rho_{t+1}|$; and $\omega_{t+1}(v) = \omega_t(v)$ for all other vertices.*
- (3) *$t := t + 1$ and go back to (1).*

One condition the cleaning model has, like the chip-firing game, but not the edge-searching problem, is that the cleaning process is to be automatic, continuing on for the lifetime of the network. That is, a final configuration of brushes (after a network has been cleaned) is to be a viable initial configuration of brushes (to clean the network again).

Definition 2.2 *The **parallel brush number**, $pb(G)$, is the minimum number of brushes needed to clean G using a parallel cleaning process.*

It was shown in [10] that $b(G) = pb(G)$ for any network G . However, as shown in Figure 2, a final configuration of brushes using the parallel cleaning process may not be a viable initial configuration. Consequently, an obvious question and the focus of this paper, is “how many extra brushes are required to continually clean a network?”

Let δ_A denote the Kronecker delta where

$$\delta_A = \begin{cases} 1 & \text{if } A \text{ is true} \\ 0 & \text{otherwise.} \end{cases}$$

Definition 2.3 Let G be a network with initial configuration $\omega_0^0 = \omega_0$. Then G can be **continually cleaned** using the parallel cleaning process beginning from configuration ω_0 if for each $s \in \mathbb{N} \cup \{0\}$, G can be cleaned in parallel using initial configuration ω_0^s , yielding the final configuration $\omega_{K_s}^s$ where $\omega_0^{s+1} = \omega_{K_s}^s$.

Given initial configuration ω_0^s , let D_t^s denote the set of dirty vertices at step t ; let $D_t^s(v) = |N(v) \cap D_t^s| \delta_{v \in D_t^s}$ denote the number of dirty edges incident to v at step t ; and let $\rho_{t+1}^s = \{v \in D_t^s : \omega_t^s(v) \geq D_t^s(v)\}$ denote the set of vertices that may be cleaned at step t .

Definition 2.4 The **continual parallel brush number**, $cpb(G)$, of a network G is the minimum number of brushes needed to continually clean G using a parallel cleaning process.

3 Continual One-Step Cleaning

In this section, we address the problem of continually cleaning all the edges of a network in one step.

Definition 3.1 Let ω_0 be an initial configuration that will continually clean network G . For each $s \in \mathbb{N} \cup \{0\}$, if $D_1^s(v) = 0$ for each $v \in V(G)$, then we say that G can be **continually one-step cleaned**.

Note that the condition “ $D_1^s(v) = 0$ for each $v \in V(G)$ ” in the previous definition requires only that each edge is cleaned in one step, not each vertex (which would require the condition “ $D_1^s = \emptyset$ ”).

Given configuration ω_0 , suppose network G can be continually one-step cleaned. As each edge must be cleaned after the first step in each cleaning, the set of vertices ρ_1^s for $s \in \mathbb{N} \cup \{0\}$ must be a vertex cover.

Theorem 3.2 Given initial configuration ω_0 , G can be continually one-step cleaned if and only if ρ_1^s is a vertex cover for every $s \in \mathbb{N} \cup \{0\}$.

For every network $G = (V, E)$ there is an initial configuration ω_0 such that the network can be continually one-step cleaned. Simply set $\omega_0(v) = \deg(v)$ for every $v \in V$. Then at $t = 0$, every vertex is cleaned (using a total of $2|E|$ brushes). Thus, the following definition is natural.

Definition 3.3 The **continual one-step brush number** of network G , denoted by $cpb_1(G)$, is the minimum number of brushes needed to continually one-step clean G .

It is also not difficult to see that $cpb_1(G) \geq |E|$ as every edge of G must be traversed by at least one brush in one step. It will be shown in Theorem 3.4 that for a network $G = (V, E)$, $|E|$ and $2|E|$ are the only two possible values for $cpb_1(G)$.

The chip firing game (see [4] for example), begins after each vertex is assigned a (finite) number of chips. At each step, one vertex, with at least as many chips as its degree, is fired (whereupon it sends one chip to each neighbour). In the parallel chip firing game, introduced in [3], at each step, every vertex with at least as many chips as its degree is fired. Note that the process of one-step cleaning a network is simply a case of the parallel chip-firing game. It was observed in [3] that for the parallel chip-firing game, any chip configuration converges in a finite time T to a limit cycle of period p . As the one-step cleaning process is an instance of the parallel chip-firing game, the observation can be applied here also and in the proof of Theorem 3.4, the focus is on the steps from T to $T + p - 1$.

Theorem 3.4 *For any connected network $G = (V, E)$,*

$$cpb_1(G) = \begin{cases} |E| & \text{if } G \text{ is bipartite} \\ 2|E| & \text{otherwise.} \end{cases}$$

Proof: Suppose first that G is bipartite with partite sets: V_0 and V_1 . For every $v \in V$, let

$$\omega_0^0(v) = \deg(v)\delta_{v \in V_0}.$$

Then for every $v \in V$,

$$\omega_1^0(v) = \deg(v)\delta_{v \in V_1}$$

and $\omega_2^0(v) = \omega_1^0(v)$ with $K_0 = 2$. It is easy to see that $\omega_t^s = \omega_t^{s \bmod 2}$ for all $t \in \{0, 1, 2\}$ and all $s \in \mathbb{N} \cup \{0\}$. Thus $cpb_1(G) \leq |E|$ and, since $cpb_1(G) \geq |E|$ for every network G , the assertion follows.

Suppose now that G is not bipartite. Let $\omega_0^0 = \omega_0$ be an initial configuration that will continually one-step clean G .

Recall that as the continual one-step cleaning of a network G is an instance of the parallel chip-firing game, it must eventually become cyclic with period p (since there are finitely many possible configurations of brushes); that is, $\omega_0^T = \omega_0^{T+p}$ for some $T, p \in \mathbb{N}$. Thus, $\omega_0^{T+t} = \omega_0^{T+(t \bmod p)}$ for $t \geq 0$ where p is the length of the period.

During a period of length p , suppose a vertex v is cleaned k_v times in the first step. Then in one period, $k_v \deg(v)$ brushes traverse from v to its neighbours and $\sum_{u \in N(v)} k_u$ brushes traverse from neighbours of v to v . So

$$k_v \deg(v) = \sum_{u \in N(v)} k_u \tag{1}$$

for every $v \in V$. It will be shown that all vertices are cleaned in the first step the same number of times during the period of length p .

For a contradiction, suppose that there is a vertex v such that $k_v \geq k_u$ for $u \in N(v)$ and there is a vertex $w \in N(v)$ such that $k_v > k_w$. But, using (1), we get the contradiction

$$\deg(v) = \sum_{u \in N(v)} \frac{k_u}{k_v} < \deg(v).$$

Thus, all vertices must be cleaned in the first step exactly k times during each period.

It is clear that $k \leq p$. Note also that if a vertex is not cleaned at the first step of the i^{th} cleaning process, it will be cleaned at the first step of the $i + 1^{\text{th}}$ cleaning process. So $p/2 \leq k \leq p$. In fact, k is greater than $p/2$. For a contradiction, suppose that $k = p/2$: every vertex is cleaned in the first step during the i^{th} process for $i = 2, 4, 6, \dots$ (or for $i = 1, 3, 5, \dots$) only and the vertex set can be decomposed into two sets ρ_0^i and ρ_1^i . Note that no two neighbours are cleaned in the first step of the i^{th} process (for any i); if they are, then they both are not cleaned in the first step of the $i + 1^{\text{th}}$ process and the edge joining them is not cleaned in the first step of the $i + 1^{\text{th}}$ process. However, this implies there is no edge between sets ρ_0^i and ρ_1^i . This is a contradiction because G is not bipartite.

Now, let us fix an edge $uv \in E$. Since both u and v are cleaned more than $p/2$ times during the period of length p , there must be some step $T_{u,v}$ at which both vertices are cleaned. Then a brush b_u is sent from u to v and a brush b_v is sent from v to u . We will show that we can force these two brushes to stay in vertices u or v from that step of the process on.

Let $A(t)$ denote the event that exactly one of b_u, b_v is at each of u, v ; $A(T_{u,v})$ trivially holds. If both u and v are cleaned at step t and $A(t)$ holds, then brushes exchange each other and $A(t + 1)$ holds as well. If only one vertex, say u , is cleaned at step t and $A(t)$ holds, then v contains at least $\deg(v) + 1$ brushes at time $t + 1$, including brushes b_u and b_v ; v is cleaned at time $t + 1$, and since we have enough brushes we can leave b_u at v and move b_v to u , thus $A(t + 2)$ holds as well. So brushes b_u, b_v stay at u or v at every time $t \geq T_{u,v}$ and every edge is associated with two brushes for $t \geq \max\{T_{u,v} : uv \in E\}$. Finally, the total number of brushes is at least $2|E|$, and since $cpb_1(G) \leq 2|E|$ for every network G , the proof is complete. ■

The concept can easily be extended to continual k -step cleaning of a network G where $cpb_k(G)$ is the minimum number of brushes needed to continually k -step clean a network G .

4 Continual Parallel Cleaning

In this section, the focus shifts from examining the number of brushes needed to continually clean a network in one step, to the more general problem of determining the minimum number of brushes, $cpb(G)$, needed to continually clean a network G (regardless of the number of steps). Section 4.1 determines the continual brush number for cycles and Section 4.2 determines the continual brush number for trees. Although the continual brush number for complete bipartite networks is determined in Section 4.3, only upper and lower bounds will be determined for complete networks. Theorem 4.7 and Corollary 4.11 bound the continual brush number for a complete network K_n between $5/16n^2 + O(n)$ and $4/9n^2 + O(n)$.

Given initial configuration $\omega_0 = \omega_0^0$, suppose that a network G can be continually cleaned. As G is finite, there are finitely many possible configurations of brushes. Thus, the cleanings must eventually become cyclic with period p , that is, settle into some periodic sequence of initial conditions. That is, $\omega_0^T = \omega_0^{T+p}$ for some $T, p \in \mathbb{N}$. Thus, $\omega_0^{T+t} = \omega_0^{T+(t \bmod p)}$ for $t \geq 0$ where p is the length of the period.

Let ω_0 be an initial configuration for a network G . The problem of determining $cpb(G)$ is also made difficult by the fact that G may be cleaned k times (for some value of k), but not $k + 1$ times. For example, given the network G with initial configuration ω_0^0 shown in Figure 3 (a), the network can be cleaned, leaving final configuration ω_7^0 . Using $\omega_7^0 = \omega_0^1$ as the initial configuration for the second cleaning as in Figure 3 (b), G is cleaned leaving

final configuration ω_5^1 . Using $\omega_5^1 = \omega_0^2$ as the initial configuration for the third cleaning as in Figure 3 (c), G cannot be cleaned a third time.




Figure 3: A network G with several ‘good’ initial configurations.

4.1 Continual Parallel Cleaning: Cycles

Given a cycle C_3 with vertices a, b, c , it is easy to see that $b(C_3) = 2$ and $cpb(C_3) = 3$. Suppose two brushes are initially placed at vertex a . Then in the first step, a is cleaned, sending one brush to each of b, c . In the second step, both b and c are cleaned, each sending a brush to the other. In the final configuration, each of b, c has one brush. Certainly, this is not a viable initial configuration to clean C_3 again. Thus, at least 3 brushes are needed to clean C_3 and it is easy to see this is an exact value.

Theorem 4.1 *For any cycle C_n with $n \geq 2$,*

$$cpb(C_n) = \begin{cases} 2 & \text{if } n \text{ is even} \\ 3 & \text{if } n = 3 \\ 4 & \text{otherwise.} \end{cases}$$

Proof: Let $C_n = (V, E)$ be a cycle on n vertices and label the vertices of C_n such that $v_i v_{i+1} \in E$ for all $i \in \{0, 1, \dots, n-2\}$ and $v_0 v_{n-1} \in E$. It is clear that at least two brushes are needed to clean a cycle. If $n = 3$, then by the argument above, $cpb(C_3) = 3$. Suppose n is even. If two brushes are placed at v_0 in the initial configuration, then after the network has been cleaned (in parallel), the two brushes must be located at $v_{n/2}$. As the network is symmetric, the final configuration is equivalent to the initial configuration. Thus, $cpb(C_n) = 2$ if n is even.

Suppose n is odd and $n > 3$. Two brushes are placed at v_0 in the initial configuration. Then in the final configuration, there is one brush located at $v_{(n-1)/2}$ and one at $v_{(n+1)/2}$. As C_n cannot be cleaned a second time, $cpb(C_n) > 2$.

Now we try to clean C_n using 3 brushes. Initially place two brushes at v_0 . The third brush must be located at either $v_{(n-1)/2}$ or $v_{(n+1)/2}$ in order to be able to clean the network again. By symmetry, suppose it is located at $v_{(n-1)/2}$. Once C_n has been cleaned, there are

two brushes at $v_{(n-1)/2}$ and one brush at $v_{(n+1)/2}$. The network can be cleaned a second time, but note that the brush located at $v_{(n+1)/2}$ has no impact on the second cleaning process. That is, the final configuration (of the second cleaning) yields one brush at each of $v_{(n+1)/2}$, v_0 and v_{n-1} . As C_n cannot be cleaned a third time, $cpb(C_n) > 3$.

It is easy to see that $cpb(C_n) \leq 4$ by initially placing two brushes at each of v_0, v_1 . After the first cleaning there are two brushes at $v_{(n+1)/2}$ and one at each of v_0, v_1 . After the second cleaning, there are two brushes at each of v_0, v_1 , which is where they began originally. ■

4.2 Continual Parallel Cleaning: Trees

Theorem 4.2 *Let $G = (V, E)$ be a network that contains a bridge $ab \in E$. Suppose G is cleaned by a parallel cleaning process using $b(G) = pb(G)$ brushes. Then the vertices a and b are not cleaned at the same time step.*

Proof: Let $\mathfrak{P} = \{(\omega_t, D_t)\}_{t=0}^K$ be a parallel cleaning process cleaning G with $b(G)$ brushes. For a contradiction, suppose that a and b are cleaned at time step k , $1 \leq k \leq K$. To show that the number of brushes used by \mathfrak{P} is not minimum, we distinguish two cases.

First, suppose there is at least one brush at vertex a or b in the initial configuration. Without loss of generality, assume that $\omega_0(b) > 0$ (the case $\omega_0(a) > 0$ is symmetric). Let $\rho = (\rho_1, \rho_2, \dots, \rho_K)$ be the parallel cleaning sequence corresponding to \mathfrak{P} where ρ_i is the set of vertices that are cleaned (in parallel) at time step i . To show that \mathfrak{P} does not use a minimum number of brushes, let us define a new initial configuration τ_0 that uses less brushes to sequentially clean G :

$$\tau_0(v) = \begin{cases} \omega_0(v) - 1 & \text{if } v = b \\ \omega_0(v) & \text{otherwise.} \end{cases}$$

The following sequential cleaning sequence \mathfrak{C} cleans G starting from the initial configuration τ_0 . First, clean (in any order) the set of vertices ρ_1 , then ρ_2 , and so on, up to (and including), ρ_{k-1} , then $\rho_k \setminus \{b\}$, then vertex b , then ρ_{k+1} , ρ_{k+2} , and so on. Note that for each vertex v , the set of vertices cleaned before v in \mathfrak{P} is contained in the set of vertices cleaned before v in \mathfrak{C} . Vertex b can be cleaned as it received an extra brush from vertex a in an earlier step, and all the other vertices can be cleaned as $\tau_0 = \omega_0$ for every vertex except b .

Now, suppose that $\omega_0(a) = \omega_0(b) = 0$. Let (A, B) be a partition of V such that $a \in A$, $b \in B$ and ab is the only edge adjacent to a vertex of A and a vertex of B in G . Let $G_A = G[A \cup \{b\}]$ and $G_B = G[B \cup \{a\}]$ where $G[S]$ denotes the subnetwork of G induced by the set $S \subseteq V$. As \mathfrak{P} is a parallel cleaning sequence where no brush that started from a vertex in A cleans an edge in $G[B]$, no brush that started from a vertex in B cleans an edge in $G[A]$ and the edge ab is cleaned by both a brush that started from a vertex in A and a brush that started from a vertex in B , there exist sequential cleaning sequences for G_A and G_B with initial configurations τ_0^A and τ_0^B where $\tau_{|B|+1}^B(a) = 1$, $\tau_{|A|+1}^A(b) = 1$, and $\tau_0^A + \tau_0^B = \omega_0$. Therefore the following initial configuration τ_0 for G uses one brush less than ω_0 :

$$\tau_0(v) = \begin{cases} \tau_0^A(v) & \text{if } v \in A \\ \tau_{|B|+1}^B(v) & \text{otherwise.} \end{cases}$$

It is straightforward to compose the cleaning sequence for G_A (starting from τ_0^A) and the reversed cleaning sequence for G_B (starting from $\tau_{|B|+1}^B$) into a cleaning sequence for G starting from τ_0 , as the only edge cleaned by both sequences is ab and it is cleaned in the same direction. Thus, indeed \mathfrak{P} does not use a minimum number of brushes which gives a contradiction and finishes the proof. \blacksquare

As every edge of a tree is a bridge, the following corollary follows.

Corollary 4.3 *Consider a parallel cleaning sequence cleaning a tree T using $b(T) = pb(T)$ brushes. The set of vertices cleaned at each time step is an independent set.*

Theorem 4.4 *For any tree T , $cpb(T) = b(T) = pb(T)$.*

Proof: Consider a parallel cleaning process \mathfrak{P} cleaning tree T with $b(T)$ brushes. It is clear that we can also clean T using the following sequential cleaning \mathfrak{C} : clean (in any order) the set of vertices cleaned at the first step of \mathfrak{P} , then the set of vertices cleaned at the second step, and so on. Note also that it follows from Corollary 4.3 that the final configuration of \mathfrak{C} is the same as in \mathfrak{P} .

It was shown in [10] that the *sequential* cleaning process is reversible. Given an initial configuration ω_0 , let ω_n be the final configuration of brushes after a network G has been sequentially cleaned. It was shown in [10] that using initial configuration $\tau_0 = \omega_n$, G can be sequentially cleaned, yielding final configuration $\tau_n = \omega_0$.

Consequently \mathfrak{C} is reversible which implies that \mathfrak{P} is reversible. \blacksquare

4.3 Continual Parallel Cleaning: Complete Bipartite networks

Theorem 4.5 *For any complete bipartite network $K_{m,n}$, $cpb(K_{m,n}) = \lceil mn/2 \rceil$.*

Proof: It was shown in [9] that $b(K_{m,n}) \geq \lceil mn/2 \rceil$; consequently $cpb(K_{m,n}) \geq pb(K_{m,n}) \geq b(K_{m,n}) \geq \lceil mn/2 \rceil$.

Let M and N be the partite sets of $K_{m,n}$, where $M = \{u_1, u_2, \dots, u_m\}$ and $N = \{v_1, v_2, \dots, v_n\}$. The proof is broken into two cases: first, assuming m is even and second, assuming both m, n are odd.

Suppose m is even and set

$$\omega_0^0(v) = \begin{cases} n & \text{if } v = u_i \text{ for } i = 1, 2, \dots, m/2 \\ 0 & \text{otherwise.} \end{cases}$$

In step 1, the first $m/2$ of the u_i 's are cleaned and at step 2, all v_j 's are cleaned as $\omega_1^0(v_j) = m/2$ for all j . The final configuration is

$$\omega_3^0(v) = \omega_2^0(v) = \begin{cases} n & \text{if } v = u_i \text{ for } i = m/2 + 1, m/2 + 2, \dots, m \\ 0 & \text{otherwise} \end{cases}$$

and as the final configuration is equivalent to the initial one, the process is continual.

Suppose now that both m, n are odd and set

$$\omega_0^0(v) = \begin{cases} 1 & \text{if } v = u_i \text{ for } i = 1, 2, \dots, (m+1)/2 \\ m & \text{if } v = v_j \text{ for } j = 1, 2, \dots, (n-1)/2 \\ 0 & \text{otherwise.} \end{cases}$$

At step 1, the first $(n-1)/2$ of the v_j 's are cleaned as $\omega_0^0(v_j) = \deg(v_j)$.

$$\omega_1^0(v) = \begin{cases} (n+1)/2 & \text{if } v = u_i \text{ for } i = 1, 2, \dots, (m+1)/2 \\ (n-1)/2 & \text{if } v = u_i \text{ for } i = (m+1)/2 + 1, \dots, m \\ 0 & \text{otherwise.} \end{cases}$$

At step 2, there are $(n+1)/2$ dirty vertices in N , so all vertices in M with $(n+1)/2$ brushes are cleaned.

$$\omega_2^0(v) = \begin{cases} (n-1)/2 & \text{if } v = u_i \text{ for } i = (m+1)/2 + 1, \dots, m \\ (m+1)/2 & \text{if } v = v_j \text{ for } j = (n+1)/2, \dots, n \\ 0 & \text{otherwise.} \end{cases}$$

At step 3, there are $(m-1)/2$ dirty vertices in M , so all vertices in N with $(m+1)/2$ brushes are cleaned. Every edge has now been cleaned and

$$\omega_4^0(v) = \omega_3^0(v) = \begin{cases} n & \text{if } v = u_i \text{ for } i = (m+1)/2 + 1, \dots, m \\ 1 & \text{if } v = v_j \text{ for } j = (n+1)/2, \dots, n \\ 0 & \text{otherwise.} \end{cases}$$

It is easy to see the network can be cleaned a second time: Let $\omega_0^1 = \omega_4^0$. Then $(m-1)/2$ vertices of M are cleaned in the first step, each sending a brush to each vertex in N . Then $(n+1)/2$ vertices in N each have $(m+1)/2$ brushes while the other vertices of N each have $(m-1)/2$. In the second step, the $(n+1)/2$ vertices of N with $(m+1)/2$ brushes are cleaned. Then the $(m+1)/2$ dirty vertices of M each have $(n+1)/2$ brushes and the $(n-1)/2$ dirty vertices of N each have $(m-1)/2$ brushes. In the third step, each dirty vertex of M is cleaned, leaving a brush at each, and each dirty vertex of N now has m brushes. Finally, $\omega_3^1 = \omega_4^1$ and as $\omega_4^1 = \omega_0^0$, the process must be continual.

In both cases, the network is cleaned with $\lceil mn/2 \rceil$ brushes. ■

Let $G = (V, E)$ be a bipartite network with initial configuration ω_0 that will parallel clean G using $pb(G) = b(G)$ brushes (note that we are considering G to be cleaned *once*). One might be tempted to suggest that if the vertices cleaned at each step form an independent set, then $cpb(G) = pb(G) = b(G)$. This is actually not the case. If the vertices cleaned at each step form an independent set, then G may be cleaned a second time, but it does not guarantee that G can be continually cleaned using the initial configuration ω_0 . This is easily demonstrated by the example in Figure 4. Given the initial configuration in Figure 4 (a), network G can be cleaned such that the vertices cleaned at each step in the process form an independent set. The vertices cleaned at each time step *mod* 2 are indicated in Figure 4 (b). However, note from the explanation of Figure 3, that the network can be cleaned a second time, but not a third time.




Figure 4: A network G with several initial configurations.

It is interesting to note, however, that given the initial configuration shown in Figure 5 (a), it is easy to see that G can be continually cleaned using $cpb(G) = b(G)$ brushes.

The final configuration for the first cleaning can be used as an initial configuration for the second cleaning, whose final configuration is the same as the initial configuration of the first round.




Figure 5: A network G with an initial configuration such that $cpb(G) = b(G)$.

This leads to the following question: does there always exist an initial configuration that yields $cpb(G) = b(G)$ for a bipartite network G ? The answer to this question is a resounding ‘no’ and an example of a bipartite network G for which $cpb(G) > b(G)$ is given in Figure 6. One can easily determine that $b(G) \leq 4$ (in fact $b(G) = 4$), by placing three brushes at x_2 and one brush at y_2 in the initial configuration. We now try to clean G in parallel using four brushes. There are only five initial configurations that can be used to clean network G from Figure 6 using four brushes:

$$\begin{aligned}
(1) \quad \omega_0(v) &= \begin{cases} 3 & \text{if } v = x_2 \\ 1 & \text{if } v = y_2 \\ 0 & \text{otherwise;} \end{cases} & (2) \quad \omega_0(v) &= \begin{cases} 3 & \text{if } v = y_2 \\ 1 & \text{if } v = x_2 \\ 0 & \text{otherwise;} \end{cases} \\
(3) \quad \omega_0(v) &= \begin{cases} 2 & \text{if } v = x_1 \\ 1 & \text{if } v = x_2 \\ 1 & \text{if } v = y_2 \\ 0 & \text{otherwise;} \end{cases} & (4) \quad \omega_0(v) &= \begin{cases} 2 & \text{if } v = y_1 \\ 1 & \text{if } v = y_2 \\ 1 & \text{if } v = x_2 \\ 0 & \text{otherwise;} \end{cases} \\
(5) \quad \omega_0(v) &= \begin{cases} 1 & \text{if } v = x_4 \\ 1 & \text{if } v = x_6 \\ 1 & \text{if } v = y_4 \\ 1 & \text{if } v = y_6 \\ 0 & \text{otherwise;} \end{cases}
\end{aligned}$$




Figure 6: A bipartite network G for which $cpb(G) > b(G)$.

If G is parallel cleaned using initial configuration (1), (2), (3), or (4), the final configuration is (5). However, if G is cleaned using initial configuration (5), the final configuration, shown in Figure 7, cannot be used to clean the network a second time. Consequently $cpb(G) > pb(G) = b(G)$ for the network shown in Figures 6 and 7. Given the initial configuration with two brushes at each of u_1, v_1 , one can easily determine that $cpb(G) = 5$. See Figure 8 for a configuration of 5 brushes that continually clean.

This leads us to the following question.

Question 4.6 For a bipartite network G , can the difference between $cpb(G)$ and $b(G)$ be arbitrarily large?




Figure 7: A bipartite network G with initial configuration (5) and final configuration indicated with circles.




Figure 8: Alternative drawing of the network in Figure 6 with 5 brushes that continually clean.

It certainly is not trivial to check all possible non-isomorphic (connected) bipartite networks on n vertices. However, we implemented and ran programs written in C/C++ to determine that $cpb(G) = b(G)$ for every bipartite network $G = (V, E)$ with $|V| \leq 11$; Figure 6 illustrates the only network on 12 vertices for which $cpb(G) \neq b(G)$ (the programs can be downloaded from [14]).

4.4 Continual Parallel Cleaning: Complete networks

In this section, we show that the continual parallel brush number for complete networks K_n is bounded between $5/16n^2 + O(n)$ and $4/9n^2 + O(n)$ by Theorems 4.7 and Corollary 4.11. Initial configurations that will continually parallel clean K_n using $4/9n^2 + O(n)$ brushes are also given in Theorems 4.8, 4.9, and 4.10.

As determined in Section 4.1, $cpb(K_3) = 3$. Similarly, one can easily determine that $cpb(K_4) = 5$. However, determining $cpb(K_n)$ for larger values of n is certainly not trivial and this section concludes with a table of calculated values of $cpb(K_n)$.

Theorem 4.7 *For any complete network K_n , $cpb(K_n) \geq \frac{5}{16}n^2 + O(n)$.*

Proof: Label the vertices of K_n as v_0, v_1, \dots, v_{n-1} and note that

$$cpb(K_n) \geq pb(K_n) = b(K_n) = \lfloor n^2/4 \rfloor.$$

It was determined in [10] that to clean K_n with $\lfloor n^2/4 \rfloor$ brushes, we must use the following assignment:

$$\omega_0^0(v_i) = \begin{cases} n - 2i - 1 & \text{if } i \leq \lfloor \frac{n-1}{2} \rfloor \\ 0 & \text{otherwise.} \end{cases}$$

Note that up to a relabelling of the vertices, none of the vertices can have less brushes than given in the assignment. It is easy to see that after K_n has been parallel cleaned once, the final configuration is

$$\omega_T^0(v_i) = \begin{cases} 0 & \text{if } i \leq \lfloor \frac{n-1}{2} \rfloor \\ \lfloor \frac{n}{2} \rfloor & \text{otherwise.} \end{cases}$$

After the first cleaning, there are $\lceil n/2 \rceil$ vertices, each with $\lfloor n/2 \rfloor$ brushes. However, if we wish to continually clean K_n , then in the final configuration (of the first cleaning), there must be at least one vertex which has at least $n - 1$ brushes, at least one additional vertex with at least $n - 3$ brushes, and so on.

More precisely, if $\lceil n/2 \rceil$ is even, then to continually parallel clean K_n , at least an additional

$$1 + 3 + 5 + \dots + \lceil n/2 \rceil - 1 = \lceil n/2 \rceil^2/4$$

brushes are required. If $\lceil n/2 \rceil$ is odd, then to continually parallel clean K_n , at least an additional

$$2 + 4 + 6 + \dots + \lceil n/2 \rceil - 1 = (\lceil n/2 \rceil^2 - 1)/4$$

brushes are required.

$$cpb(K_n) \geq \begin{cases} \frac{5}{16}n^2 & \text{if } n \text{ is even and } \lceil n/2 \rceil \text{ is even} \\ \frac{5}{16}n^2 - \frac{1}{4} & \text{if } n \text{ is even and } \lceil n/2 \rceil \text{ is odd} \\ \frac{5}{16}n^2 + \frac{1}{8}n - \frac{3}{16} & \text{if } n \text{ is odd and } \lceil n/2 \rceil \text{ is even} \\ \frac{5}{16}n^2 + \frac{1}{8}n - \frac{7}{16} & \text{if } n \text{ is odd and } \lceil n/2 \rceil \text{ is odd.} \end{cases}$$

■

We now determine upper bounds for the continual parallel brush number of K_n . Note that it is broken into three results: $n = 3k$ in Theorem 4.9; $n = 3k + 1$ in Theorem 4.10; and $n = 3k + 2$ in Theorem 4.8. We only give the proof of Theorem 4.8 since all the initial configurations, and hence the proofs, are very similar,

Theorem 4.8 *For any complete network on $n = 3k + 2$ vertices,*

$$cpb(K_{3k+2}) \leq 4k^2 + 4k + 2.$$

Proof: Let $n = 3k + 2$, label the vertices of K_n as $v_0, v_1, \dots, v_{3k+1}$, and set

$$\omega_0(v_i) = \begin{cases} k + 1 & \text{if } i = k, k + 1, \dots, 2k + 1 \\ i & \text{otherwise.} \end{cases}$$

As $\omega_0(v_{3k+1}) = \deg(v_{3k+1})$, we clean v_{3k+1} at step 1. At step 2 we clean $2^{2-1} = 2$ vertices; at step 3 we clean $2^{3-1} = 4$ vertices; at step 4 we clean 2^{4-1} vertices; and so on up to step $\lceil \log_2(k + 1) \rceil - 1$ when we clean $2^{\lceil \log_2(k+1) \rceil - 2}$ vertices. This can be done as

$$\begin{aligned} \omega_{t-1}(v_i) &= i + 2^{t-1} - 1 \\ &\geq (3k - 2^t + 3) + 2^{t-1} - 1 \\ &= 3k - 2^{t-1} + 2 \\ &= D_{t-1}(v_i) \end{aligned}$$

for each $i \in \{3k + 2 - 2^t + 1, 3k + 2 - 2^t + 2, \dots, 3k + 2 - 2^t + 2^{t-1}\}$.

For a vertex v_i that is cleaned at step $t = \lceil \log_2(k + 1) \rceil$: it had i brushes initially, it has received $2^{t-1} - 1$ brushes in the previous steps, and $D_{t-1}(v_i) = (3k + 1) - (2^{t-1} - 1)$. There are a total of $(3k + 2) - (2k + 2) - (2^{t-1} - 1) = k - (2^{t-1} - 1)$ vertices cleaned at step t , so in total, $\omega_{t+2}(v_i) = i + (2^{t-1} - 1) - 3k - 1 + (2^{t-1} - 1) + k - (2^{t-1} - 1) - 1 = i + 2^{t-1} - 2k - 3$.

At step $t + 1 = \lceil \log_2(k + 1) \rceil + 1$, $D_t(v_i) = 2k + 2$ and $\omega_t(v_i) = \omega_0(v_i) + k$ for all $v_i \in D_t$. Thus, we clean $v_k, v_{k+1}, \dots, v_{2k+1}$ at step $t + 1$.

At step $t + 2$, $\omega_{t+1}(v_i) = i + 2k + 2$ for all $v_i \in D_{t+1}$ so the remaining k vertices are cleaned, yielding final configuration of

$$\omega_{t+2}(v_i) = \begin{cases} i + 3 \cdot 2^{t^*-1} - 3k - 4 & \text{for } i = 3k - 2^{\lceil \log_2(k+1) \rceil - 1} + 3, \dots, 3k + 1 \\ i + 2^{\lceil \log_2(k+1) \rceil - 1} - 2k - 3 & \text{for } i = 2k + 2, \dots, 3k - 2^{\lceil \log_2(k+1) \rceil - 1} + 2 \\ k & \text{for } i = k, k + 1, \dots, 2k + 1 \\ i + 2k + 1 & \text{for } i = 0, 1, \dots, k - 1 \end{cases}$$

where $t^* = \lceil \log_2(3k - i + 3) \rceil$ is the step at which v_i was cleaned.

By renaming the sets and variables, we have an initial configuration $\omega_0^1 = \omega_{t+2}$ equivalent to $\omega_0 = \omega_0^0$. ■

Theorem 4.9 Let $n = 3k$ and label the vertices of K_n as $v_0, v_1, \dots, v_{3k-1}$. If

$$\omega_0(v_i) = \begin{cases} k & \text{if } i = k-1, k, \dots, 2k-1 \\ i & \text{otherwise,} \end{cases}$$

then K_{3k} can be continual parallel cleaned using this initial configuration of $4k^2 - k + 1$ brushes.

Theorem 4.10 Let $n = 3k + 1$ and label the vertices of K_n as v_0, v_1, \dots, v_{3k} . If

$$\omega_0(v_i) = \begin{cases} k & \text{if } i = k, k+1, \dots, 2k \\ i & \text{otherwise,} \end{cases}$$

then K_{3k+1} can be continual parallel cleaned using this initial configuration of $4k^2 + k$ brushes.

From Theorem 4.9, Theorem 4.10, and Theorem 4.8:

Corollary 4.11

$$cpb(K_n) \leq \begin{cases} 4/9n^2 - 1/3n + 1 & \text{if } n = 3k \\ 4/9n^2 - 5/9n + 1/9 & \text{if } n = 3k + 1 \\ 4/9n^2 - 4/9n + 10/9 & \text{if } n = 3k + 2. \end{cases}$$

So in general, $cpb(K_n) \leq 4/9n^2 + O(n)$.

We implemented and ran programs written in C/C++ to determine the values of $cpb(K_n)$ for $n = 3, 4, \dots, 20$ — see Table 1 (the programs and results can be downloaded from [14]). It seems that there should be at least one initial configuration with the maximum value of $n - 1$ that can be used to start the process that continues forever, but no proof of this fact is known. However, by checking all such configurations we were able to find an upper bound for $cpb(K_n)$ for $n = 21, 22, \dots, 26$. We also conjecture that those are, in fact, the exact values.

n	$b(K_n)$	$cpb(K_n)$	n	$b(K_n)$	$cpb(K_n)$	n	$b(K_n)$	$cpb(K_n)$
3	2	3	11	30	50	19	90	150
4	4	5	12	36	60	20	100	170
5	6	10	13	42	66	21	110	≤ 173
6	9	15	14	49	76	22	121	< 197
7	12	18	15	56	95	23	132	≤ 214
8	16	25	16	64	105	24	144	≤ 243
9	20	33	17	72	122	25	156	≤ 258
10	25	39	18	81	135	26	169	≤ 279

Table 1: $b(G)$ and $cpb(K_n)$ for $n = 3, 4, \dots, 26$.

Figure 9 plots $b(K_n)/cpb(K_n)$ for $n = 3, 4, \dots, 20$ and it remains an open question to determine the value of $\lim_{n \rightarrow \infty} b(K_n)/cpb(K_n)$ if it exists. From Table 1, the upper bound of Theorems 4.10, 4.8 and 4.9 sometimes give the minimum number of brushes. Thus, we conjecture that:

$$\lim_{n \rightarrow \infty} b(K_n)/cpb(K_n) = (1/4)/(4/9) = 9/16.$$




Figure 9: A network of $b(K_n)/cpb(K_n)$ versus n (from 3 to 20).

5 Acknowledgement

This work was made possible by the facilities of

- the Shared Hierarchical Academic Research Computing Network SHARCNET, Ontario, Canada (www.sharcnet.ca): 8,082 CPUs,
- the Atlantic Computational Excellence Network ACEnet, Memorial University of Newfoundland, St. John's, NL, Canada (www.ace-net.ca): 412 CPUs.

References

- [1] B. Alspach, Searching and Sweeping Graphs: A Brief Survey, *Le Matematiche*, 2004, **59**, 5-37.
- [2] N. Alon, P. Prałat, and N. Wormald, Cleaning d -regular graphs with brushes, *SIAM Journal on Discrete Mathematics*, submitted, 19pp.
- [3] J. Bitar, E. Goles, Parallel chip firing games on graphs, *Theoretical Computer Science* **92** (1992), 291–300.
- [4] A. Björner, L. Lovász, W. Shor, Chip-firing games on graphs, *European Journal of Combinatorics* **12** (1991) 283–291.
- [5] B. Hobbs, J. Kahabka, 1995. Underwater Cleaning Technique Used for Removal of Zebra Mussels at the Fitzpatrick Nuclear Power Plant. *Proceedings of The Fifth International Zebra Mussel and Other Aquatic Nuisance Organisms Conference*, Toronto, Canada, February 1995.
- [6] S. R. Kotler, E. C. Mallen, K. M. Tamms, Robotic Removal of Zebra Mussel Accumulations in a Nuclear Power Plant Screenhouse, *Proceedings of The Fifth International Zebra Mussel and Other Aquatic Nuisance Organisms Conference*, Toronto, Canada, February 1995.

- [7] S. McKeil, Chip Firing Cleaning Processes, MSc Thesis, Dalhousie University, 2007.
- [8] N. Megiddo, S. L. Hakimi, M. R. Garey, D. S. Johnson & C. H. Papadimitriou, The complexity of searching a graph, *Journal of the Association for Computing Machinery*, 1988, **35**, 18-44.
- [9] M. E. Messinger, Methods of Decontaminating a Network, PhD Thesis, Dalhousie University, 2008.
- [10] M. E. Messinger, R. J. Nowakowski, and P. Prałat, Cleaning a Network with Brushes, *Theoretical Computer Science* **399** (2008) 191–205.
- [11] M. E. Messinger, R. J. Nowakowski, P. Prałat, and N. Wormald, Cleaning random d-regular graphs with brushes using a degree-greedy algorithm, *Proceedings of the 4th Workshop on Combinatorial and Algorithmic Aspects of Networking (CAAN2007)*, Lecture Notes in Computer Science, Springer, 2007, 13–26.
- [12] M. E. Messinger, R. J. Nowakowski, The Robot Cleans Up, *Proceedings of the 2nd Annual International Conference on Combinatorial Optimization and Applications (COCOA '08)*, Lecture Notes in Computer Science, Springer, 2008, 309–318.
- [13] P. Prałat, Cleaning random graphs with brushes, *Australasian Journal of Combinatorics*, submitted, 16pp.
- [14] P. Prałat, programs written in C/C++,
<http://www.mathstat.dal.ca/~pralat/index.php?page=publications>.