Longest Induced Cycles in Circulant Graphs

Elena D. Fuchs

Department of Mathematics
University of California, Berkeley, Berkeley, CA
lenfuchs@berkeley.edu

Submitted: Aug 19, 2004; Accepted: Jun 13, 2005; Published: Oct 13, 2005
Mathematics Subject Classifications: 05C88, 05C89

Abstract

In this paper we study the length of the longest induced cycle in the unit circulant graph $X_n = Cay(\mathbb{Z}_n; \mathbb{Z}_n^*)$, where $\mathbb{Z}_n^*$ is the group of units in $\mathbb{Z}_n$. Using residues modulo the primes dividing n, we introduce a representation of the vertices that reduces the problem to a purely combinatorial question of comparing strings of symbols. This representation allows us to prove that the multiplicity of each prime dividing n, and even the value of each prime (if sufficiently large) has no effect on the length of the longest induced cycle in X_n. We also see that if n has r distinct prime divisors, X_n always contains an induced cycle of length $2^r + 2$, improving the $r \ln r$ lower bound of Berrezbeitia and Giudici. Moreover, we extend our results for X_n to conjunctions of complete k_i-partite graphs, where k_i need not be finite, and also to unit circulant graphs on any quotient of a Dedekind domain.

1 Introduction

For a positive integer n, let the unit circulant graph $X_n = Cay(\mathbb{Z}_n, \mathbb{Z}_n^*)$ be defined as follows:

(1) The vertex set of X_n, denoted by $V(n)$, is $\mathbb{Z}_n$, the ring of integers modulo n.

(2) The edge set of X_n is denoted by $E(n)$, and, for $x, y \in V(n)$, $\{x, y\} \in E(n)$ if and only if $x - y \in \mathbb{Z}_n^*$, where $\mathbb{Z}_n^*$ is the set of units in the ring $\mathbb{Z}_n$.

The central problem adressed in this paper is to find the length of the longest induced cycle in X_n. This problem was first considered by Berrizbeitia and Giudici [1], who were motivated by its applications to chromatic uniqueness.

Throughout the paper, we let $n = p_1^{a_1} p_2^{a_2} \dots p_r^{a_r}$, where the p_i are distinct primes, and $a_i \geq 1$. Then we denote the length of the longest induced cycle in X_n by $m(n)$. We let $M(r) = \max_n m(n)$, where the maximum is taken over all n with r distinct prime divisors. In [1], Berrizbeitia and Giudici bound $M(r)$ by

$$r \ln r \leq M(r) \leq 9r!.$$

A simple change to the proof of the upper bound provided in [1] yields the better upper bound of $M(r) \le 6r!$.

Our goal is to determine better bounds for $M(r)$, as well as to extend what we find to other graphs. In Section 2, we introduce a useful representation of the vertices in X_n according to their residues modulo the prime divisors of n. This representation immediately yields several helpful properties of the longest induced cycles in these graphs. In particular, we prove that we can disregard the multiplicities of the prime divisors of n, so we can reduce our problem to square-free n. Also, we show that $m(n)$ depends only on r, and in fact $m(n) = M(r)$ as long as the primes dividing n are all large enough. In Section 3, we use the vertex representation introduced in Section 2 to construct an induced cycle of length $2^r + 2$ in the graph X_n, where n has r distinct prime divisors, thus raising the lower bound on $M(r)$ substantially. We also note that this construction is valid for any n, no matter what its prime divisors are, so this provides a lower bound for $m(n)$. Section 4 contains a generalization of our results to conjunctions of complete k_i-partite graphs, as well as to unit circulant graphs on products of local rings, which include the unit circulant graphs on Dedekind rings.

2 Residue Representation

Recall that $n = p_1^{a_1} p_2^{a_2} \cdots p_r^{a_r}$, where the p_i are prime. We will represent the vertices of X_n in a way that will reduce the process of finding induced cycles in X_n to checking for similarities between strings of numbers in an array.

It is clear that the following is equivalent to the definition of $E(n)$ in the introduction:

Observation 2.1. For $x, y \in V(n)$, we have that $\{x, y\} \in E(n)$ if and only if

$$x \not\equiv y \pmod{p_i}, \text{ for all } 1 \le i \le r.$$

Likewise, $\{x, y\} \notin E(n)$ if and only if

$$x \equiv y \pmod{p_i}, \text{ for some } 1 \le i \le r.$$

So, in fact, to know whether x and y are adjacent we need only their residues modulo the primes p_i. With this in mind, we introduce the following representation of the vertices:

Definition 2.2.

(i) Let $x \in V(n)$, such that

$$x \equiv \alpha_i \pmod{p_i}, \text{ where } 1 \le i \le r \text{ and } 0 \le \alpha_i < p_i.$$

We then define the *residue representation* of x to be the unique string $\alpha_1 \alpha_2 \cdots \alpha_r$, where α_k is the *kth term*, and we write $x \approx \alpha_1 \alpha_2 \cdots \alpha_r$.

(ii) Let $x, y \in V(n)$. If the kth term of the residue representation of x is the same as the kth term of the residue representation of y, we say that x has a *coincidence* with y.

Combining Observation 2.1 and Definition 2.2, vertices $x, y \in V(n)$ are adjacent if and only if x has no coicidences with y. So, in fact, the only property of the residues modulo p_i that we use in constructing induced cycles is that they form a set of size p_i, and we verify that a subgraph is an induced cycle by checking that consecutive vertices do not have any coincidences, and that any pair of non-consecutive vertices has at least one coincidence.

Also, we note that for n not square-free, a string may be the residue representation of multiple vertices. For example, if $n = 12$, both 0 and 6 have residue representation 00. However, the adjacency of vertices depends only on their residue representations, and, by the Chinese Remainder Theorem, every string represents at least one vertex.

This representation greatly simplifies inspection of induced cycles. In fact, we can extend residue representation for a vertex to any induced subgraph:

Definition 2.3.

(i) Let S be an induced subgraph of X_n, where $V(S) = (v_0, v_2, \ldots, v_{k-1})$, with $v_i \approx \alpha_{i1}\alpha_{i2}\cdots\alpha_{ir}$, and $0 \le i \le k-1$. We then define the *residue representation of* S to be the array

$$\begin{array}{cccc} \alpha_{01} & \alpha_{02} & \cdots & \alpha_{0r} \\ \alpha_{11} & \alpha_{12} & \cdots & \alpha_{1r} \\ \vdots & \vdots & & \vdots \\ \alpha_{(k-1)1} & \alpha_{(k-1)2} & \cdots & \alpha_{(k-1)r}. \end{array}$$

(ii) The *residue set* of S is the set of residues

$$\{\alpha_{ij} \mid 0 \le i \le k-1, 1 \le j \le r\}$$

used in its residue representation.

So, if an induced subgraph S is a k-cycle in X_n, we can permute the rows of the residue representation of S so that the ith row has a coincidence with the jth row if and only if $i - j \not\equiv \pm 1 \pmod{k}$. Figure 1 displays the residue representation of an induced 6-cycle for $r = 2$ and for $r = 3$.

An important property of an induced cycle of length greater than 4 is that it cannot contain two vertices with the same residue representation.

Proposition 2.4. *The residue representation of a k-cycle C, with $k > 4$, cannot contain two identical rows.*

Proof. Suppose there are two vertices x and y in C that have the same residue representation. Then a vertex z of C has no coincidence with x if and only if it has no coincidence with y, meaning that x and y have precisely the same neighbors in C. However, a vertex in an induced cycle is adjacent to exactly two other vertices in the cycle, so C can have at most 4 vertices, contradicting $k > 4$. Thus the residue representation of C cannot contain two identical rows. ☐

$$\begin{array}{cc} 0 & 0 \\ 1 & 1 \\ 0 & 2 \\ 1 & 0 \\ 0 & 1 \\ 1 & 2 \end{array} \qquad \begin{array}{ccc} 0 & 0 & 0 \\ 1 & 1 & 1 \\ 0 & 0 & 2 \\ 1 & 2 & 0 \\ 0 & 0 & 1 \\ 1 & 1 & 2 \end{array}$$

Figure 1: In these residue representations of an induced 6-cycle for $r = 2$ on the left, and for $r = 3$ on the right, it is easy to see that two consecutive rows (including the 1st and 6th rows) have no coincidences, and any two non-consecutive rows have at least one coincidence. The residue set for each cycle is $\{0, 1, 2\}$.

It is important that, once we have written an induced cycle in terms of its residue representation, we can permute the residues in each column to obtain an induced cycle of equal length.

Observation 2.5. Let the jth column in the residue representation of an induced k-cycle C in X_n be

$$\begin{array}{c} \alpha_{0j} \\ \alpha_{1j} \\ \vdots \\ \alpha_{(k-1)j}\,, \end{array}$$

and suppose this column contains l_j distinct residues, $\{s_1, s_2, \ldots, s_{l_j}\}$. Then let π be a permutation of $\{s_1, s_2, \ldots, s_{l_j}\}$, and replace the jth column of C by

$$\begin{array}{c} \pi(\alpha_{0j}) \\ \pi(\alpha_{1j}) \\ \vdots \\ \pi(\alpha_{(k-1)j}). \end{array}$$

We then have a new induced k-cycle in X_n, since we have not changed the coincidences between any of the rows in C.

We now use the Observation 2.5 to define isomorphisms between induced k-cycles in X_n.

Definition 2.6. Two induced k-cycles, C and C', are called *isomorphic* if, for every j, the jth column of the residue representation of C' is obtained by permuting the residues in the jth column of C, as described in Observation 2.5.

Note that the first two rows in Figure 1 are 000 and 111. Because of this, all of the rows that are not adjacent to either of the first two have to contain both a 0 and a 1. Similarly, the third row in the cycle must contain a 0, and the last row in the cycle must contain a 1. This is a useful criterion for induced cycles in general.

Remark 2.7. Any induced cycle C in X_n is isomorphic to an induced cycle C' of the same length so that the first two rows in the residue representation of C' are $00\cdots0$ and $11\cdots1$.

In order to obtain such a C', we need only to map the first two elements in every column of C to 0 and 1, respectively. Note that the first two elements in each column are always different – if they were not, the first and the second row in the residue representation of C would have a coincidence, which contradicts their adjacency.

This tells us that all but four of the rows in our induced cycles will have to contain both a 0 and a 1, which may limit the residue sets and consequently the lengths of the cycles.

Another interesting fact that becomes evident with the use of residue representation is the following proposition.

Proposition 2.8. *The value $M(r)$ increases with r. Specifically, if X_n contains an induced cycle of length k, and $q > 2$ is a prime not dividing n, then X_{qn} also contains a cycle of length k. If k is even, we can also allow $q = 2$.*

Proof. Let $n = p_1^{a_1} p_2^{a_2} \cdots p_r^{a_r}$, where the exponents a_i are positive integers, and p_i are distinct primes. Suppose X_n contains an induced cycle C of length k. We denote the residue representations of the vertices of C by $v_0, v_1, \ldots, v_{k-1}$, where each v_i is a string of length r. Let $n' = qn$, where $q \neq 2$ is prime, $q \neq p_i$ for all $1 \leq i \leq r$. Then we will show that $X_{n'}$ also contains a cycle of length k by constructing an induced cycle C' in $X_{n'}$, denoting the residue representations of the vertices of C' by $w_0, w_1, \ldots, w_{k-1}$.

If k is even, let $w_i = v_i0$ for even i, and let $w_i = v_i1$ for odd i. Notice that we do not introduce any coincidences between two rows that were adjacent in C, so two consecutive rows in C' are adjacent, as desired. Similarly, if $\{v_i, v_j\} \notin E(n)$, they have a coincidence, say, in the lth term. Then w_i and w_j have a coincidence in the lth term, and so $\{w_i, w_j\} \notin E(n')$. Thus we introduce no new adjacencies in the construction of C', so C' is indeed an induced k-cycle in $X_{n'}$.

If k is odd, let $w_i = v_i1$ for odd i, let $w_i = v_i0$ for even $i \neq k-1$, and let $w_{k-1} = v_{k-1}2$ (this is possible since $q \neq 2$). Again, we note that we do not introduce any coincidences between two rows that were adjacent in C, so two consecutive rows in C' are adjacent, as desired. Also, if $\{v_i, v_j\} \notin E(n)$, we have that $\{w_i, w_j\} \notin E(n')$ by the argument above. Thus we introduce no new adjacencies in the construction of C', so C' is indeed an induced k-cycle in $X_{n'}$.

By starting with a cycle C in $X(n)$ that has length $M(r)$, we see that $M(r+1) \geq M(r)$, as desired. □

Corollary 2.9. *If $r \geq 2$, and n is square-free, then $m(n) \geq 6$.*

Proof. For $r = 2$, we have constructed a 2-cycle of length 6 in Figure 1, so $M(2) \geq 6$. Proposition 2.8 shows that $M(r)$ is nondecreasing, so we have that, if $r > 2$, $M(r) \geq M(2) \geq 6$, as desired. □

We now prove that, in calculating $m(n)$, we need consider only those n that are square-free.

Theorem 2.10. *For $n = p_1^{a_1} p_2^{a_2} \cdots p_r^{a_r}$, and $n' = p_1 p_2 \cdots p_r$, where $r \neq 1$, $m(n) = M(n')$.*

Proof. (1) First we show that $m(n) \geq M(n')$. In particular, we show X_n contains cycles of length $M(n')$. Note that since n and n' have the same prime divisors, if $x, y < n$, then $x - y \in \mathbb{Z}_n^*$ if and only if $x - y \in \mathbb{Z}_{n'}^*$. So, in particular, the induced subgraph of X_n on vertices $0, 1, \ldots, n' - 1$ is precisely $X_{n'}$. Thus any induced cycle on $X_{n'}$ can be mapped to an induced cycle in $\{0, 1, \ldots, n' - 1\} \subset X_n$, and so there is an induced cycle of length $M(n')$ in X_n, as desired.

(2) Now we show that $m(n) \leq M(n')$, or that there is no induced cycle of length greater than $M(n')$ in X_n. Since n' is square-free, Corollary 2.9 implies that $M(n') \geq 6$. Suppose there is an induced cycle, C_l, of length $l > M(n')$ in X_n. Then, in particular, $l > 6$. Using residue representation, write C_l in terms of residues $\pmod{p_1, p_2, ..., p_r}$. If no two vertices in C_l are denoted by the same string of residues, then we can view the residue representation of C_l as a residue representation of an induced l-cycle in $X_{n'}$. Since $l > M(n')$, this contradicts the assumption that $M(n')$ is the maximum length of an induced cycle in $X_{n'}$. Thus there exist two vertices in C_l that have identical residue representations. However, by Proposition 2.4, this means $l \leq 4$, contradicting the previous deduction that that $l > 6$. We conclude that, indeed, there are no induced cycles of length $l > M(n')$ in X_n. □

Proposition 2.11. *Let $n' = p$, and $n = p^a$ where p is a prime and $a > 1$. Then $M(n') = 3$, and $m(n) = 4$. So, $M(1) = 4$.*

Proof. Since the only non-unit in $\mathbb{Z}_p$ is 0, $X_{n'}$ is a complete graph on p vertices, and the longest induced cycle in $X_{n'}$ must hence have length 3. From Part (2) of the proof of Theorem 2.10, we deduce that $m(n) \leq 4$. In fact, $m(n) = 4$, since the subgraph $(0, 1, p, p + 1)$ is an induced cycle in X_n. □

Proposition 2.12. *For $n = p_1^{a_1} p_2^{a_2} \cdots p_r^{a_r}$ where the p_i are large, $m(n) = M(r)$.*

Proof. Let n' be a positive integer with exactly r prime divisors, such that $m(n') = M(r)$, and let C be a longest induced cycle in $X_{n'}$. Assume each p_i is larger than the number of residues that appear in the residue representation of C. Then there is a subgraph S of X_n, such that the residue representation of S is the same as the residue representation of C. So S is an induced cycle of length $M(r)$. Hence $m(n) = M(r)$. Thus, as long as the prime divisors of n yield enough residues for a residue representation of the longest cycle in $X_{n'}$, where $M(n') = M(r)$, we will have $m(n) = M(r)$. □

3 A Lower Bound on $M(r)$

One important asset of introducing residue representation is that it gives us a way to construct a good lower bound on $M(r)$; we achieve the following lower bound as our main result in this section.

Theorem 3.1. *For all positive integers n with $r > 1$ distinct prime divisors, we have $M(r) \geq 2^r + 2$.*

In this section, we construct an induced subgraph of X_n with 2^r+2 vertices, where r is the number of distinct prime divisors of n, and provide two specific cycles produced by this construction. We will then prove that this subgraph is indeed a cycle, and thus show that Theorem 3.1 holds.

In order to construct an induced 2^r+2-cycle in X_n, where $n = p_1 p_2 \cdots p_r$, we first introduce some definitions, which are discussed in detail in [4], p. 433.

(i) An *n-bit Gray Code* is an ordered, cyclic sequence of the 2^n n-bit binary strings called *codewords*, such that successive codewords differ by the complementation of a single bit, and the starting codeword is taken to be $(00\cdots 0)$. We write this sequence in the form of a matrix, as shown below.

(ii) A *Reflective Gray Code* (RGC) is defined recursively as follows: A 1-bit RGC is merely the 2×1 matrix $\binom{0}{1}$. If an r-bit RGC is the $2^r \times r$ binary matrix

$$\begin{pmatrix} G_0 \\ G_1 \\ \vdots \\ G_{2^r-1}, \end{pmatrix}$$

then we define the $(r+1)$-bit RGC to be the $2^{r+1}\times(r+1)$ binary matrix

$$\begin{pmatrix} 0G_0 \\ 0G_1 \\ 0G_2 \\ \vdots \\ 0G_{2^r-1} \\ 1G_{2^r-1} \\ 1G_{2^r-2} \\ \vdots \\ 1G_1 \\ 1G_0 \end{pmatrix}.$$

Henceforth, we fix r and index the codewords by $0, 1, \ldots, 2^r-1 \pmod{2^r}$, denoting the ith codeword in an r-bit RGC by G_i, and the ith codeword in a k-bit RGC, where $k \neq r$, by $G_i^{(k)}$.

(iii) The *flip bit* in the jth codeword of a RGC is the position of the one bit that has changed from the $(j-1)$st codeword.

We will construct an induced subgraph of X_n whose residue representation consists of the rows $v_0, v_1, \ldots, v_M$, where $M = 2^r+1$, and $\{v_i, v_j\} \in E$ if and only if $i-j \equiv \pm 1 \pmod{2^r+2}$. Let $v_{M-1} \approx 0100\cdots 0$, and $v_M \approx 122\cdots 2$. We define the rows $\{v_i : i \text{ even}, i \neq M-1\}$ by using the first half of an r-bit RGC with a slight modification. Let $\widehat{G}_i$, for $i \neq 0$ be the ith codeword G_i in an r-bit RGC, with the flip bit replaced by a 2. Let $\widehat{G}_0 = G_0$. Then we define the even-indexed rows as follows: $v_{2i} = \widehat{G}_i$, for $0 \leq i < 2^{r-1}$.

0	0	0
1	1	1
0	0	2
1	1	0
0	2	1
1	0	0
0	1	2
1	0	1
0	1	0
1	2	2

0	0	0	0
1	1	1	1
0	0	0	2
1	1	1	0
0	0	2	1
1	1	0	0
0	0	1	2
1	1	0	1
0	2	1	0
1	0	0	1
0	1	1	2
1	0	0	0
0	1	2	1
1	0	1	0
0	1	0	2
1	0	1	1
0	1	0	0
1	2	2	2

Figure 2: We construct two cycles using residue representation and our lower bound construction. On the left is an induced 10-cycle for the graph X_n, where n has three prime divisors ($r = 3$). On the right is an induced 18-cycle for the graph X_n, where n has four prime divisors ($r = 4$). Note that the rows in both cycles are derived as described from a 3-bit Reflective Gray Code and a 4-bit Reflective Gray Code, respectively.

We define the odd-indexed rows as follows: for $0 \leq i \leq 2^{r-1}$, let $v_{2i+1} = \overline{G_i}$, the complement of G_i. So the subgraph we have constructed is

$$\{\widehat{G}_0, \overline{G}_0, \widehat{G}_1, \ldots, \widehat{G}_{2^{r-1}-1}, \overline{G}_{2^{r-1}-1}, v_{M-1}, v_M\}.$$

This gives us a subgraph consisting of $(2^r + 2)$ vertices.

In Figure 2, we display this construction for $r = 3$ and $r = 4$.

To prove Theorem 3.1, we must show that the subgraph we have constructed is indeed an induced cycle. This can be reduced to showing that the following properties hold.

(i) Vertex v_k is adjacent to v_l if $k - l \equiv \pm 1 \pmod{2^r + 2}$. In other words, $\{v_0, v_1, \ldots v_M\}$ is a cycle.

(ii) If neither k nor l equals $M - 1$ or M, and $|k - l| > 1$, then v_k is not adjacent to v_l.

(iii) Vertex v_M is not adjacent to v_l for $i \neq 0, M - 1$, and vertex v_{M-1} is not adjacent to v_l for $i \neq M - 2, M$.

Proof of Theorem 3.1.

(i) First we show that any two consecutive rows among $v_0, v_1, \ldots, v_{M-2}$ correspond to adjacent vertices. Among these rows, no odd-indexed row contains a 2, and an even-indexed row v_{2i} is merely the complement of v_{2i+1} with one bit replaced by a 2. Thus every odd-indexed row among $v_0, v_1, \ldots, v_{M-2}$ has no coincidences with the row immediately above it. Also, since any two consecutive codewords G_i and G_{i+1} in an r-bit RGC differ only in the flip bit of G_{i+1}, the codeword $\overline{G}_i$ differs from G_{i+1} everywhere except in the flip bit. However, in modifying G_i to $\widehat{G}_i$ for $0 \le i < 2^{r-1}$, we have replaced every flip bit by a 2, so $v_{2i+1} = \overline{G}_i$, (which will contain no 2's), will differ completely from $v_{2i+2} = \widehat{G}_{i+1}$ if $i \ne 2^{r-1} - 1$. Thus every odd-indexed row among $v_0, v_1, \ldots, v_{M-4}$ is adjacent to the row immediately below it.

It remains to show that v_M is adjacent to v_{M-1}, that v_M is adjacent to v_0 (these two claims are trivial by inspection), and that v_{M-2} is adjacent to v_{M-1}. Note that v_{M-1} is precisely $G_{2^{r-1}-1}$, since, by definition,

$$G_{2^{r-1}-1} = 0G^{(r-1)}_{2^{r-2}-1} = 01G^{(r-2)}_0 = 0100\cdots0.$$

Also, v_{M-2} is, by definition, $\overline{G}_{2^{r-1}-1}$. Thus, indeed, v_{M-2} is adjacent to v_{M-1}, and we have that $\{v_0, v_1, \ldots v_M\}$ is a cycle.

(ii) It is trivial to show that no two rows whose indices have the same parity are adjacent, since all even-indexed rows begin with a 0 and are thus not adjacent to each other, while all odd-indexed rows begin with a 1 and are also not adjacent to each other.

Now, take an even-indexed row v_{2i}, with $0 \le i < 2^{r-1}$, and an odd-indexed row v_{2j+1}, with $0 \le j < 2^{r-1}$, such that $i \ne j$ and $i \ne j+1$. Suppose for the sake of contradiction that v_{2i} is adjacent to v_{2j+1}.

By definition, $v_{2j+1} = \overline{G}_j$, $v_{2i} = \widehat{G}_i$, and $i \ne j$ by assumption. By the definition of a RGC, G_j and G_i differ in at least one bit. Since $i - j \not\equiv 1 \pmod{2^r}$, then G_j and G_i must differ in a bit that is not a flip bit for G_i. Therefore $v_{2j+1} = \overline{G}_j$ will have at least one coincidence with $v_{2i} = \widehat{G}_i$, and so v_{2i} and v_{2j+1} are not adjacent, contrary to our supposition.

So, indeed, if neither k nor l equals $M-1$ or M, and $|k-l| > 1$, then v_k is not adjacent to v_l.

(iii) Since v_M begins with a 1, it is not adjacent to any of the odd-indexed rows, which also all begin with a 1. Similarly, because all of the even-indexed rows except v_0 and v_{M-1} have a 2 in some spot after the initial 0, and will thus have a coincidence with $v_M \approx 122\cdots2$, no even-indexed row except v_0 and v_{M-1} will be adjacent to v_M.

Since v_{M-1} begins with a 0, it is not adjacent to any of the even-indexed rows, which all begin with a 0 as well. Also, note that $v_{M-2} = v_{2^r-1} = \overline{G}_{2^{r-1}-1} = 1011\cdots1$ is the complement of v_{M-1}, and that all odd-indexed rows except v_M are distinct and contain only 0's and 1's. Thus all odd-indexed rows except v_M either complement or have a coincidence with $V_{M-1} = 0100\cdots0$. So all odd-indexed rows except for v_{M-2} and v_M are not adjacent to v_{M-1}.

Thus we have that vertex v_M is not adjacent to v_i for $i \ne 0, M-1$, and vertex v_{M-1} is not adjacent to v_i for $i \ne M-2, M$. $\square$

Note that, for any $n = p_1p_2\cdots p_r$, where $p_1 < p_2 < \cdots < p_r$ are primes, the cycle constructed above does not depend on the choice of p_i. The first column of the cycle's residue representation contains residues 0 and 1 only, allowing for $p_1 = 2$, and the residue set of the cycle is $\{0, 1, 2\}$, which puts no bounds on the rest of the primes p_i.

Also, Theorem 2.10 implies that our construction of a $(2^r + 2)$-cycle for $n' = p_1 p_2 ... p_r$, $r > 1$ holds for $n = {p_1}^{a_1} {p_2}^{a_2} ... {p_r}^{a_r}$, while Proposition 2.11 implies that the lower bound in Theorem 3.1 holds for $r = 1$.

4 Generalizing to Other Graphs

A natural question to ask is what properties of the circulant graph X_n are necessary to obtain the results we have. It is noted in [1] that, for p prime and a a positive integer, X_{p^a} is complete p-partite. In fact, this tells us that for $n = p_1^{a_1} p_2^{a_2} \cdots p_r^{a_r}$, X_n is the conjunction $X_{p_1^{a_1}} \wedge X_{p_2^{a_2}} \wedge \cdots \wedge X_{p_r^{a_r}}$ of graphs $X_{p_1^{a_1}}, X_{p_2^{a_2}}, \cdots, X_{p_r^{a_r}}$, where a conjunction of graphs is defined as follows:

Definition 4.1. Let the graph G_1 have vertex set $V(G_1)$ and edge set $E(G_1)$, and graph G_2 have vertex set $V(G_2)$ and edge set $E(G_2)$. Then the *conjunction* $G_1 \wedge G_2$ has vertex set $V(G_1 \wedge G_2) = V(G_1) \times V(G_2)$, and (v_1, v_2) is adjacent to (u_1, u_2) if $v_1 u_1 \in E(G_1)$, and $v_2 u_2 \in E(G_2)$.

Interestingly, our results can be extended to any conjunction $G_1 \wedge G_2 \wedge \cdots \wedge G_r$, where each G_i is complete k_i-partite. Let $S = \{k_1, k_2, \ldots, k_r\}$ be an r-tuple of positive integers. Let $\mathcal{G}^S = \{G | G = G_1 \wedge G_2 \wedge \cdots \wedge G_r\}$, where G_i is a complete k_i-partite graph. Denote the length of the longest induced cycle in $G \in \mathcal{G}^S$ by $\mathcal{M}(S)$, and define $\mu(r) = \max_S \mathcal{M}(S)$ to be the length of the longest induced cycle in all graphs in $\mathcal{G}^S$, where S contains r integers.

Theorem 4.2. *For $r > 1$, we have that $\mu(r) = M(r)$.*

To prove Theorem 4.2, we will create for conjunctions of k_i-partite graphs a representation similar to residue representation. Then, using this representation, we will show how cycles in $G \in \mathcal{G}^S$ and X_n are related.

Definition 4.3. Let $S = \{k_1, k_2, \ldots, k_r\}$, and let $G \in \mathcal{G}^S, G = G_1 \wedge G_2 \wedge \cdots \wedge G_r$. Label the partitions in G_i by $\{0, 1, 2, \ldots, k_i - 1\}$. Let $v = (v_1, v_2, \ldots, v_r) \in V(G)$, where v_i belongs to partition α_i in G_i. Then the *partition representation* of v is $\alpha_1 \alpha_2 \cdots \alpha_r$, and we say $v \simeq \alpha_1 \alpha_2 \cdots \alpha_r$.

We can define the partition representation of a subgraph of $G \in \mathcal{G}^S$ as we defined the residue representation of a subgraph of X_n. Namely, an induced subgraph on $\{x_1, x_2, \ldots, x_l\}$ is written as an array of partition representations of the vertices x_i. Note that an induced subgraph in G is a cycle precisely when its partition representation satisfies the conditions needed for the residue representation of an induced cycle in X_n – no two non-consecutive rows can have coincidences, and two non-consecutive rows must have at least one coincidence.

Proof of Theorem 4.2.

(1) First we show that $M(r) \geq \mu(r)$. Suppose $S = \{k_1, k_2, \ldots, k_r\}$, and $G \in \mathcal{G}^S$ contains an induced cycle C of length $\mu(r)$, whose partition representation is

$$\begin{array}{cccc} \alpha_{11} & \alpha_{12} & \cdots & \alpha_{1r} \\ \alpha_{21} & \alpha_{22} & \cdots & \alpha_{2r} \\ \vdots & \vdots & & \vdots \\ \alpha_{\mu(r)1} & \alpha_{\mu(r)2} & \cdots & \alpha_{\mu(r)r} \end{array}.$$

Note that, applying Proposition 2.4 to partition representations, no two rows above are identical if $\mu(r) > 4$. So, if $\mu(r) > 4$, let $n = p_1 p_2 \cdots p_r$, where the p_i are prime and $p_i \geq \max\{\alpha_{1i}, \alpha_{2i}, \ldots, \alpha_{\mu(r)i}\}$. Then the partition representation of C above is in fact also the residue representation of an induced cycle in X_n, and so X_n contains a cycle of length $\mu(r)$, as desired. If $\mu(r) \leq 4$, we know that $M(r) \geq \mu(r)$, since $M(1) = 4$, and $M(r)$ increases with r by Proposition 2.8.

(2) Now we show that $\mu(r) \geq M(r)$. Let X_n, where $n = p_1 p_2 \cdots p_r$, contain an induced cycle of length $M(r)$. Then $X_n \in G^{\{p_1, \ldots, p_r\}}$, so $\mu(r) \geq M(r)$, as desired. □

Since our original problem concerns the circulant graph X_n, we are also interested in circulant graphs to which our results generalize. In particular, we are interested in those graphs $G = Cay(A; A^*)$, where A is a ring, A^* is the group of units in A, and the graph G is defined as follows:

(1) The vertex set $V(G)$ of G is the set of elements in A.

(2) If $x, y \in V(G)$ then $\{x, y\} \in E(G)$, the edge set of G, if and only if $x - y \in A^*$.

We know that we can extend our results to any graph G that is a conjunction of complete k_i-partite graphs for some k_i. Note that, surprisingly, k_i need not be finite, and, in fact, our circulant graph need not contain a finite number of vertices. For this, we rely on a partition using the Chinese Remainder Theorem. One can refer to an algebra text such as [2], pp. 92-97 for the basic facts about rings and ideals needed to prove when such a partition gives us the desired graph structure.

Definition 4.4. A *local ring* is a ring that contains only one maximal ideal.

With this definition, we can show that a unitary circulant graph on a product of local rings is a conjunction of complete k_i-partite graphs.

Theorem 4.5. *Let $A_1, A_2, \ldots A_r$ be local rings, and let $\mathfrak{m}_i$ be the one maximal ideal in A_i. If $A = A_1 \times A_2 \times \cdots \times A_r$, then the circulant graph $Cay(A; A^*)$ is a conjunction of complete k_i-partite graphs, for some nonzero k_i.*

Theorem 4.5 lets us extend our results to various unitary circulant graphs. In particular, it allows us to generalize to unitary circulant graphs on Dedekind rings.

Definition 4.6. A *Dedekind domain* ([3]) is an integral domain R such that

(1) Every ideal in R is finitely generated;

(2) Every nonzero prime ideal is a maximal ideal;

(3) R is integrally closed in its field of fractions

$$K = \{\alpha/\beta : \alpha, \beta \in R, \beta \neq 0\}.$$

A Dedekind ring is simply a quotient of a Dedekind domain.

If R is a Dedekind domain, and $\mathfrak{m}_i$ is a maximal ideal of R, then $R/\mathfrak{m}_i$ is a field and thus contains only one maximal ideal, (0), and $R/\mathfrak{m}_i^{a_i}$ contains only the maximal ideal $\mathfrak{m}_i$, so $R/\mathfrak{m}_i^{a_i}$ is a local ring. This is essential for the following corollary.

Corollary 4.7. *Let R be a Dedekind domain and let $I = \mathfrak{m}_1^{a_1}\mathfrak{m}_2^{a_2}\cdots\mathfrak{m}_r^{a_r}$ be a nonzero, non-unit ideal in R, where $\mathfrak{m}_i$ are maximal ideals of R. Then the circulant graph $Cay(A; A^*)$ is a conjunction of complete k_i-partite graphs, for $k_i = \#(R/\mathfrak{m}_i)$.*

Proof. Since $\mathfrak{m}_i$ are the distinct maximal ideals, $\mathfrak{m}_i^{a_i} + \mathfrak{m}_j^{a_j} = R$ for all $1 \leq i < j \leq r$. Then the Chinese Remainder Theorem implies that

$$A = R/\mathfrak{m}_1^{a_1}\mathfrak{m}_2^{a_2}\cdots\mathfrak{m}_r^{a_r} = R/\mathfrak{m}_1^{a_1} \times R/\mathfrak{m}_2^{a_2} \times \cdots \times R/\mathfrak{m}_r^{a_r}$$

We have noted above that $R/\mathfrak{m}_1^{a_1}$ is local, and thus we have that A is a product of local rings. By Theorem 4.5, we have that the circulant graph $Cay(A; A^*)$ is a conjunction of complete k_i-partite graphs, for $k_i = \#(R/\mathfrak{m}_i)$. □

So, indeed, our theorems concerning $M(r)$ generalize to the maximum lengh of a cycle in unit circulant graphs on a Dedekind domain quotiented by an ideal with r distinct maximal factors. Dedekind domains are exactly those integral domains in which every ideal has a unique factorization into prime ideals, and thus are the rings of number theoretical interest. Some nice examples of the Dedekind rings that we have generalized to above are the Gaussian integers modulo $a + bi$, denoted by $\mathbb{Z}[i]/(a + bi)$; any quotient of the ring of algebraic integers in the pth cyclotomic field $\mathbb{Z}[\zeta_p]$, where ζ_p is a pth root of unity; and any quotient of $\mathbb{C}[x, y]/(y^2 - x^3 + x)$, the ring of regular functions on the elliptic curve $y^2 = x^3 - x$. Note that we also have generalized to unit circulant graphs on quotients of principal rings.

5 Acknowledgements

This research was done at the University of Minnesota Duluth. I would like to give special thanks to Joseph A. Gallian for his encouragement and support, and I would like to thank Philip Matchett, Melanie Wood, and Justin Sinz for their insightful comments. I would also like to thank Geir Helleloid for helpful conversations. Funding was provided by grants form the NSF (DMS-0137611) and NSA (H-98230-04-1-0050).

References

[1] P. Berrizbeitia, R. E. Giudici, On cycles in the sequence of unitary Cayley graphs, Discrete Math 282 (2004), 1-3.

[2] S. Lang. Algebra, 3rd edition, Springer-Verlag, NY, 2002.

[3] D. A. Marcus. Number Fields, Springer Verlag, NY, 1977.

[4] E. Reingold, J. Nievergelt, N. Deo. Combinatorial Algorithms, Prentice Hall, NJ, 1977.