A New Approach to the Periodicity Lemma on Strings with Holes ☆,☆☆

W. F. Smyth*

Algorithms Research Group, Department of Computing & Software, McMaster University, Hamilton, Ontario, Canada, L8S 4K1

Digital Ecosystems & Business Intelligence Institute, Curtin University of Technology, Perth WA 6845, Australia

Shu Wang

Algorithms Research Group, Department of Computing & Software, McMaster University, Hamilton, Ontario, Canada, L8S 4K1

Abstract

We first give an elementary proof of the periodicity lemma for strings containing one hole (variously called a "wild card" or a "don't-care" or an "indeterminate letter" in the literature). The proof is modelled on Euclid's algorithm for the greatest common divisor and is simpler than the original proof given in [BB99]. We then study the two hole case, where our result agrees with the one given in [BSH02] but is more easily proved and enables us to identify a maximum-length prefix or suffix of the string to which the periodicity lemma does apply. Finally we extend our result to three or more holes using elementary methods and state a version of the periodicity lemma that applies to all strings with or without holes. We describe an algorithm that, given the locations of the holes in a string, computes maximum length substrings to which the periodicity lemma applies, in time proportional to the number of holes. Our approach is quite different from the one in [BSH02, BS04] and also simpler.

Key words: periodicity, periodicity lemma, indeterminate string, hole

☆Supported in part by grants from the Natural Sciences & Engineering Research Council of Canada.

☆☆The authors express their gratitude to three anonymous referees, whose comments have materially improved the quality of this paper.

*Corresponding author

Email addresses: smyth@mcmaster.ca (W. F. Smyth), shuw@mcmaster.ca (Shu Wang)

1. Introduction

Over the last few years researchers have shown interest [BB99, IMM+03, BSH02] in strings that may contain ***don't-care*** letters; that is, letters $*$ that match every letter in a given alphabet Σ. More generally, several papers [HS03, HSW06, HSW08] have studied "indeterminate" strings that may contain "indeterminate" letters — those that match various subsets of Σ. In this article we study the more general model.

Let Σ be an alphabet and let λ_i, $|\lambda_i| \geq 2$, $1 \leq i \leq m$, be pairwise distinct subsets of Σ. We form a new alphabet $\Sigma' = \Sigma \cup \{\lambda_1, \lambda_2, .., \lambda_m\}$ and define a new relation ***match*** ($\approx$) on Σ' as follows:

- for every $\mu_1, \mu_2 \in \Sigma$, $\mu_1 \approx \mu_2$ if and only if $\mu_1 = \mu_2$;
- for every $\mu \in \Sigma$ and every $\lambda \in \Sigma' - \Sigma$, $\mu \approx \lambda$ and $\lambda \approx \mu$ if and only if $\mu \in \lambda$;
- for every $\lambda_i, \lambda_j \in \Sigma' - \Sigma$, $\lambda_i \approx \lambda_j$ if and only if $\lambda_i \cap \lambda_j \neq \emptyset$.

This idea seems to have first been mentioned in [FP74].

We observe that *match* is reflexive and symmetric but not necessarily transitive; for example, if $\lambda = \{a, b\}$, then $a \approx \lambda$ and $b \approx \lambda$ does not imply $a \approx b$. In this paper $\boldsymbol{x} = \boldsymbol{x}[1..n]$ is always a nonempty string on Σ' that may therefore contain some $\lambda \in \Sigma' - \Sigma$ at some position $h \in 1..n$; that is, $\boldsymbol{x}[h] = \lambda$. We refer to an occurrence of λ in $\boldsymbol{x}$ as a ***hole***, generalizing the usage in [BB99, BSH02, BS04], where always $\Sigma' = \Sigma \cup \{\Sigma\}$. Here a hole is equivalent to an ***indeterminate letter*** as defined in [HS03]. We also sometimes refer to the position h itself as a hole.

A string $\boldsymbol{x}$ has ***period*** (***strong period***) p if and only if for every $i, j \in 1..n$ such that $i \equiv j \bmod p$, $\boldsymbol{x}[i] \approx \boldsymbol{x}[j]$; $\boldsymbol{x}$ has ***weak period*** p if and only if for every $i, j \in 1..n$ such that $j = i+p$, $\boldsymbol{x}[i] \approx \boldsymbol{x}[j]$. For example, in the following table $\boldsymbol{x}$ has a weak period but not a strong period of length 2.

	1	2	3	4	5	6
$\boldsymbol{x} =$	a	b	a	$*$	a	c

On strings without holes, periodicity and weak periodicity are equivalent.

2. Strings With One Hole

We first consider strings with exactly one hole. In [BB99] a variant of the periodicity lemma [FW65] for such strings was stated, proved, and shown to be sharp:

Lemma 1. *If $\boldsymbol{x}$ with one hole has weak periods p and $q > p$, and $n \geq p+q$, then $\boldsymbol{x}$ has strong period $d = \gcd(p, q)$.*

We prove this lemma here based on the Euclidean algorithm, extending the proof given in [Smy03] for the original periodicity lemma. As observed in [BB99], it suffices to establish the case $n = p+q$, since therefore for larger n, the lemma holds for every factor of length $p+q$, hence for $\boldsymbol{x}$ itself. We first prove a preliminary result:

Lemma 2. *Suppose $\boldsymbol{x} = \boldsymbol{x}[1..p+q]$ has weak periods p and $q > p$ with a single hole $\boldsymbol{x}[h] = \lambda$.*

(a) $h \in 1..q \;\Rightarrow\; \boldsymbol{x}[1..q]$ has weak periods p and $q-p$;

(b) $h \in p+1..p+q \;\Rightarrow\; \boldsymbol{x}[p+1..p+q]$ has weak periods p and $q-p$.

Proof. We prove (a); the proof of (b) is analogous. Since $\boldsymbol{x}$ has weak periods p and $q > p$, therefore $\boldsymbol{x}[1..q]$ has weak period p. Since for $i > p$, $i+(q-p) > q$, we need consider only $i \in 1..p$. For these values of i, it follows from weak q periodicity that $\boldsymbol{x}[i] \approx \boldsymbol{x}[i+q]$ and from weak p periodicity that $\boldsymbol{x}[i+q] \approx \boldsymbol{x}[i+q-p]$. Since $h \leq q$, we know that $\boldsymbol{x}[i+q] \neq \lambda$, hence that $\boldsymbol{x}[i] \approx \boldsymbol{x}[i+q-p]$. Therefore $\boldsymbol{x}[1..q]$ also has weak period $q-p$, as required. □

Since h satisfies the hypothesis of either Lemma 2(a) or Lemma 2(b) (or both), we can always reduce $\boldsymbol{x}$ with a single hole, whose length $p+q$ is the sum of its distinct weak periods p and q, to a substring $\boldsymbol{y}$ with a single hole whose length q is the sum of its (not necessarily distinct) weak periods p and $q-p$: $\boldsymbol{y}$ is either a prefix $\boldsymbol{x}[1..q]$ or a suffix $\boldsymbol{x}[p+1..p+q]$ of $\boldsymbol{x}$. If $p = q-p$, we have computed $p = \gcd(p, q) = q/2$; if not, we can perform another reduction. Let us write $\boldsymbol{x^{(0)}} = \boldsymbol{x}$ and for $r \geq 0$, let $\boldsymbol{x^{(r+1)}}$ be the reduction (hence a substring) of $\boldsymbol{x^{(r)}}$. By the correctness of the Euclidean algorithm, a finite number $k \geq 1$ of reductions yields a string $\boldsymbol{x^{(k)}} = \boldsymbol{x^{(k)}}[1..2d]$ that contains one hole and has weak period $d = \gcd(p, q)$. But then, since $\boldsymbol{x^{(k)}}$ takes the form $\boldsymbol{uu}$, where $\boldsymbol{u} = \boldsymbol{x}[1..d]$, it actually has strong period d. We illustrate this reduction process with an example in Tables 1–4. Starting with a string $\boldsymbol{x^{(0)}}$ that has weak periods $q^{(0)} = 8$ and $p^{(0)} = 6$, we recursively reduce it to $\boldsymbol{x^{(3)}}$ that has a strong period 2.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
$\boldsymbol{x^{(0)}} =$	a	b	a	b	a	b	$*$	b	$(a$	b	a	b	a	$b)$

Table 1: $|x^{(0)}| = 14, q^{(0)} = 8, p^{(0)} = 6, q^{(0)} - p^{(0)} = 2$

Lemma 3. *If for some $r \in 1..k$, $\boldsymbol{x}^{(r)}$ has strong period d, then $\boldsymbol{x}^{(r-1)}$ also has strong period d.*

Proof. According to the nature of a reduction, $\boldsymbol{x}^{(r-1)}$ has weak periods p and $q > p$ that are divisible by $d = q-p$, and $|\boldsymbol{x}^{(r-1)}| = p+q$. We want to prove that for every $i, j \in 1..p+q$ such that $i \equiv j \bmod d$, $\boldsymbol{x}^{(r-1)}[i] \approx \boldsymbol{x}^{(r-1)}[j]$. We consider three cases:

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
$\boldsymbol{x^{(1)}} =$	$(a$	$b)$	a	b	a	b	$*$	b						

Table 2: $|x^{(1)}| = 8, q^{(1)} = 6, p^{(1)} = 2, q^{(1)} - p^{(1)} = 4$

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
$\boldsymbol{x^{(2)}} =$			$(a$	$b)$	a	b	$*$	b						

Table 3: $|x^{(2)}| = 6, q^{(2)} = 4, p^{(2)} = 2, q^{(2)} - p^{(2)} = 2$

1. both i and j lie in $x^{(r)}$;
2. one position (say i) lies in $x^{(r)}$, but not j;
3. neither i nor j lies in $x^{(r)}$.

Case (1) is straightforward since $x^{(r)}$ is strongly d periodic.

In case (2), assume without loss of generality that $\boldsymbol{x}^{(r)} = \boldsymbol{x}^{(r-1)}[1..q]$ — the proof for suffix $\boldsymbol{x}^{(r)} = \boldsymbol{x}^{(r-1)}[p+1..p+q]$ is analogous. By the weak periodicity of $\boldsymbol{x}^{(r-1)}$, $\boldsymbol{x}^{(r-1)}[j-q] \approx \boldsymbol{x}^{(r-1)}[j]$ and $\boldsymbol{x}^{(r-1)}[j-p] \approx \boldsymbol{x}^{(r-1)}[j]$, where $j-q < j-p \leq q$, so that both $j-q$ and $j-p$ are positions in $\boldsymbol{x}^{(r)}$. Since there is exactly one hole in $\boldsymbol{x}^{(r)}$, we may denote by j^* any one of $j-q, j-p$ that is *not* a hole. Since $i \equiv j \bmod d$ and d divides both p and q, $i \equiv j^* \bmod d$. Then by the strong d periodicity of $\boldsymbol{x}^{(r)}$,

$$\boldsymbol{x}^{(r-1)}[i] \approx \boldsymbol{x}^{(r-1)}[j^*] \approx \boldsymbol{x}^{(r-1)}[j].$$

Since j^* is not a hole, $\boldsymbol{x}^{(r-1)}[i] \approx \boldsymbol{x}^{(r-1)}[j]$, as required.

In case (3) we again need only consider prefix $\boldsymbol{x}^{(r)} = \boldsymbol{x}^{(r-1)}[1..q]$. Using the same argument as in case (2), we can find $j^* < q$, not a hole, such that $\boldsymbol{x}^{(r-1)}[j^*] \approx \boldsymbol{x}^{(r-1)}[j]$. But now the same construction applies also to $i > q$, allowing us to find $i^* < q$, not a hole, such that $\boldsymbol{x}^{(r-1)}[i^*] \approx \boldsymbol{x}^{(r-1)}[i]$. Since $i \equiv j \bmod d$, it follows that $i^* \equiv j^* \bmod d$, so that by the strong d periodicity of $\boldsymbol{x}^{(r)}$, $\boldsymbol{x}^{(r-1)}[i^*] \approx \boldsymbol{x}^{(r-1)}[j^*]$. Thus $\boldsymbol{x}^{(r-1)}[i] \approx \boldsymbol{x}^{(r-1)}[j]$. (In fact, in this case, $\boldsymbol{x}^{(r-1)}[i] = \boldsymbol{x}^{(r-1)}[j]$.) □

Lemma 3 allows us to reconstruct $\boldsymbol{x}$ by reversing the reduction, and shows that every intermediate substring $\boldsymbol{x^{(r)}}$ has the same strong period. Using again the example in Tables 4–1, we see that starting with $\boldsymbol{x^{(3)}}$ of strong period 2, every intermediate substring $\boldsymbol{x^{(2)}}$, $\boldsymbol{x^{(1)}}$, and eventually $\boldsymbol{x^{(0)}}$ will have the same strong period 2.

Therefore, Lemmas 2–3 imply Lemma 1, the periodicity lemma for strings with one hole.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14
$\boldsymbol{x^{(3)}} =$					a	b	$*$	b						

Table 4: $|x^{(3)}| = 4, q^{(3)} = 2, p^{(3)} = 2, q^{(3)} - p^{(3)} = 0$

3. Strings With Two Holes

Let $\boldsymbol{x} = \boldsymbol{x}[1..n]$ be a string with two holes that is weakly p, q periodic with $q > p$, where $n \geq 2(p+q)-d$, $d = \gcd(p, q)$. Let $L_0 = p+q-d, L_1 = p+q$, and observe that $L_1 > L_0 \geq q$. Consider the prefix $\boldsymbol{x_1} = \boldsymbol{x}[1..L_0]$ of length L_0 and the suffix $\boldsymbol{x_2} = \boldsymbol{x}[n-L_1+1..n]$ of length L_1. Since there are only two holes, no matter where they lie at least one of $\boldsymbol{x_1}$ and $\boldsymbol{x_2}$ must, by the periodicity lemmas for no-hole and one-hole strings, be d periodic. Of course the same statement holds for $\boldsymbol{x_1} = \boldsymbol{x}[1..L_1]$ and $\boldsymbol{x_2} = \boldsymbol{x}[n-L_0+1..n]$.

Since part of $\boldsymbol{x}$ is strongly d periodic, we are encouraged to investigate whether there is a way to extend the d periodic portion(s), perhaps to all of $\boldsymbol{x}$. The following definition provides one basis for such an extension:

Definition 4. Suppose that $\boldsymbol{x} = \boldsymbol{x}[1..n]$ is a string with at most two holes that is weakly p, q periodic, $q > p$. For $i \in L_0+1..n$, we say that $\boldsymbol{x}[1..i-1]$ is ***right-extendible*** (RE) if at least one of the following conditions holds:

1. $\boldsymbol{x}[i-p] \in \Sigma$;
2. $\boldsymbol{x}[i-q] \in \Sigma$;
3. $i+p \leq n$ and $\boldsymbol{x}[i+p-q] \in \Sigma$;

For example, in Table 5, x has weak periods $q = 6$ and $p = 4$. Since $d = gcd(6, 4) = 2$, $L_0 = 6+4-2 = 8$ and $L_1 = 6 + 4 = 10$. There is no hole in $x[1..L_0]$, therefore according to the original periodicity lemma, $x[1..L_0]$ is (strongly) d periodic. Furthermore, according to Definition 4, for all $i \in 9..13$, $x[1..i]$ is right-extendible, while $x[1..14]$ is not right-extendible.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18
$\boldsymbol{x} =$	a	b	a	b	a	b	a	b	$*$	b	$*$	b	a	b	c	b	a	b.

Table 5: Example: Right extendibility of a string with two holes

We remark that if neither condition (1) nor (2) in Definition 4 is satisfied, then both $i-p$ and $i-q$ are holes; since $\boldsymbol{x}$ contains at most two holes, therefore for $i+p \leq n$, $\boldsymbol{x}[i+p] \in \Sigma$, and so condition (3) can fail to hold only in the case that $q = 2p$ — thus $i+p-q = i-p$. This is the "special" case described in [BSH02].

We shall see in the next section that for strings with an arbitrary number of holes, a weaker (and more general) definition of RE suffices. Based on the

RE property, the following lemma allows us to extend a d periodic prefix to the right:

Lemma 5. *Suppose that a string $\boldsymbol{x}$ on Σ' with at most two holes is weakly p, q periodic, $q > p$, and let $d = \gcd(p, q)$. If $\boldsymbol{x}[1..i-1]$ is d periodic and RE, then $\boldsymbol{x}[1..i]$ is d periodic.*

Proof. We need only prove that for every $j \in 1..i$ such that $j \equiv i \bmod d, x[j] \approx x[i]$.

Suppose condition (1) of Definition 4 holds. By d periodicity, for every $j \in 1..i-1$ such that $j \equiv (i-p) \bmod d$, $\boldsymbol{x}[j] \approx \boldsymbol{x}[i-p]$. By weak p periodicity we know that $\boldsymbol{x}[i] \approx \boldsymbol{x}[i-p]$. Because $\boldsymbol{x}[i-p]$ is not a hole, it follows that for every $j \in 1..i$ such that $j \equiv i \equiv (i-p) \bmod d$, $\boldsymbol{x}[j] \approx \boldsymbol{x}[i]$, so that $\boldsymbol{x}[1..i]$ is d periodic.

The proof for condition (2) is analogous.

Suppose that neither condition (1) or condition (2) holds, but that (3) is true. By d periodicity, for every $j \in 1..i-1$ such that $j \equiv (i+p-q) \bmod d$, $\boldsymbol{x}[j] \approx \boldsymbol{x}[i+p-q]$. Since there are at most two holes, $\boldsymbol{x}[i+p] \in \Sigma$ and so $\boldsymbol{x}[i] = \boldsymbol{x}[i+p]$; by weak q periodicity, $\boldsymbol{x}[i+p] \approx \boldsymbol{x}[i+p-q]$; since moreover $\boldsymbol{x}[i+p-q] \in \Sigma$, in fact $\boldsymbol{x}[i] = \boldsymbol{x}[i+p-q]$. It follows that for every $j \in 1..i$ such that $j \equiv i \equiv (i+p-q) \bmod d$, $\boldsymbol{x}[j] \approx \boldsymbol{x}[i]$, so that again $\boldsymbol{x}[1..i]$ is d periodic. □

A symmetrical definition and lemma enable us to extend a d periodic suffix to the left:

Definition 6. Suppose that $\boldsymbol{x} = \boldsymbol{x}[1..n]$ is a string with zero or more holes that is weakly p, q periodic, $q > p$. For $i \in 1..n-L_0$, we say that $\boldsymbol{x}[i+1..n]$ is ***left-extendible*** (LE) if at least one of the following conditions holds:

1. $\boldsymbol{x}[i+p] \in \Sigma$;
2. $\boldsymbol{x}[i+q] \in \Sigma$;
3. $i > p$ and $\boldsymbol{x}[i-p+q] \in \Sigma$;

Lemma 7. *Suppose that a string $\boldsymbol{x}$ on Σ' with at most two holes is weakly p, q periodic, $q > p$, and let $d = \gcd(p, q)$. If $\boldsymbol{x}[i+1..n]$ is d periodic and LE, then $\boldsymbol{x}[i..n]$ is d periodic.* □

We see that under specified conditions, we can extend a strongly d periodic prefix/suffix of $\boldsymbol{x}$ by one to the right/left, respectively. If this process can be iterated to cover all of $\boldsymbol{x}$, then $\boldsymbol{x}$ is d periodic. We summarize our results as

Lemma 8. *Suppose that $\boldsymbol{x} = \boldsymbol{x}[1..n]$ is a string with two holes and weak periods p and $q > p$, where $n \geq L_0+L_1$, $d = \gcd(p, q)$. Then:*

(*a*) *At least one of $\boldsymbol{x}[1..L_0]$ and $\boldsymbol{x}[n-L_1+1..n]$ is d periodic.*

(*b*) *If $\boldsymbol{x}[1..L_0]$ is d periodic and for every $i \in L_0+1..n$, $\boldsymbol{x}[1..i-1]$ is RE, then $\boldsymbol{x}$ is d periodic.*

(c) *If* $\boldsymbol{x}[n-L_1+1..n]$ *is* d *periodic and for every* $i \in 1..n-L_1$, $\boldsymbol{x}[i+1..n]$ *is LE, then* $\boldsymbol{x}$ *is* d *periodic.* □

As suggested earlier, this result can also be stated in terms of $\boldsymbol{x}[1..L_1]$ and $\boldsymbol{x}[n-L_0+1..n]$; note also that it applies to strings with any form of hole, not only don't-cares. Lemma 8 basically agrees with the result given in [BSH02], where d periodicity of $\boldsymbol{x}$ is shown to depend on $\boldsymbol{x}$ being "not $(2, p, q)$-special". However, the iterative approach given here is simpler and leads directly to a straightforward $\Theta(n)$-time algorithm to compute the maximum-length d periodic suffix/prefix of $\boldsymbol{x}[1..n]$ with two holes.

To understand this better, again we consider the weakly $4, 6$ periodic two-hole string of Table 5. By Lemma 5 the 2 periodic prefix $\boldsymbol{x}[1..8]$ can be iteratively extended to the right, yielding the conclusion that $\boldsymbol{x}[1..14]$ is 2 periodic. Since none of the conditions (1)-(4) of Definition 4 is satisfied in position 15, no further extension is possible. This makes sense since $\boldsymbol{x}[15] = c$, so that $\boldsymbol{x}[1..15]$ is not 2 periodic. Observe however that even if we transform $\boldsymbol{x}$ into $\boldsymbol{x'}$ by changing position 15 from c to a, $\boldsymbol{x'}[1..14]$ can still not be right-extended, because of the definition. Nevertheless $\boldsymbol{x'}$ is in fact 2 periodic.

In order to resolve such situations, we state a more precise version of Lemma 8, as follows:

Corollary 9. *Suppose that* $\boldsymbol{x} = \boldsymbol{x}[1..n]$ *is a string with two holes* h_1 *and* $h_2 > h_1$ *and weak periods* p *and* $q > p$, *where* $n \geq L_0+L_1$, $d = \gcd(p, q)$.

(a) *If* $h_2-h_1 \neq q-p$, *then* $\boldsymbol{x}$ *is* d*-periodic.*

(b) *If* $h_2-h_1 = q-p$, *then*

  (i) $h_2+p > n$ *or* $h_1 \leq p \Rightarrow \boldsymbol{x}$ *is* d *periodic;*

  (ii) *otherwise,* $\boldsymbol{x}[h_2+p] = \boldsymbol{x}[h_1-p] \Leftrightarrow \boldsymbol{x}$ *is* d *periodic.*

Proof.

(a) If the gap between the holes is never $q-p$, then either condition (1) or condition (2) of both Definitions 4 and 6 will hold for every i. Thus one of Lemmas 5 and 7 can be used to extend the d periodic segment of $\boldsymbol{x}$ to the full range $1..n$.

(b) Suppose then that the gap between holes is exactly $q-p$. Even so, if $h_2+p > n$ (respectively, $h_1 \leq p$), there can exist no i such that conditions (1)–(3) of Definition 4 (respectively, 6) all fail to hold. Again, the d periodic segement can be extended, either right or left, to the full range.

Suppose then that $h_2+p \leq n$ *and* $h_1 > p$. Since $n \geq L_0+L_1$, either $\boldsymbol{x}[1..h_2+p-1]$ or $\boldsymbol{x}[h_1-p+1..n]$ is d periodic. In both cases, to establish whether the d periodic range can be extended (to $\boldsymbol{x}[1..h_2+p]$ or to $\boldsymbol{x}[h_1-p..n]$), it suffices to perform the single comparison

$$\boldsymbol{x}[h_2+p] \ : \ \boldsymbol{x}[h_1-p],$$

where, since two holes are accounted for, both must be regular letters in Σ. If unequal, then the d periodic range cannot be extended; if equal, then since the remainder of the string contains no holes, the entire string is d periodic.

□

This result yields the following simple constant-time algorithm:

function d-`range`$(\boldsymbol{x}, n, p, q, h_1, h_2)$
if $h_2 - h_1 \neq q - p$ **or** $h_2 + p > n$ **or** $h_1 \leq p$ **then**
    return $1, n$
elsif $\boldsymbol{x}[h_2 + p] = \boldsymbol{x}[h_1 - p]$ **then**
    return $1, n$
elsif $h_1 + h_2 > n$ **then**
    return $1, h_2 + p - 1$
else
    return $h_1 - p + 1, n$

Figure 1: For weakly p, q periodic $\boldsymbol{x}[1..n]$, $q > p$, $n \geq L_0 + L_1$, identify the maximum d periodic range that contains holes h_1 and $h_2 > h_1$.

Our methodology extends easily and naturally to three or more holes, as discussed in the next section.

4. Strings With Zero or More Holes

For a string $\boldsymbol{x}$ with three holes and length $n \geq 2L_1$, again we consider a prefix $\boldsymbol{x_1} = \boldsymbol{x}[1..L_1]$ and a suffix $\boldsymbol{x_2} = \boldsymbol{x}[n - L_1 + 1..n]$: now both of them have length L_1. Note that since there are only three holes, at least one of these substrings has no more than one hole. If at least two holes lie in $\boldsymbol{x_1}$, so that at most one hole lies in $\boldsymbol{x_2}$, then by Lemma 1 we know that $\boldsymbol{x}_2$ is d periodic; otherwise $\boldsymbol{x_1}$ is d periodic. In either case, at least a substring (prefix or suffix) of $\boldsymbol{x}$ is d periodic. Figure 2 shows possible positions of these three holes, where in this case x_1 is d periodic.

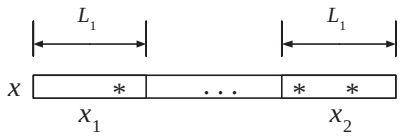

Figure 2: Possible positions of three holes

We can extend this result to any number of holes. For $d = \gcd(p, q)$, in addition to $L_0 = p + q - d, L_1 = p + q$, for $k \geq 2$ define $L_k = L_{k-2} + L_1$. Thus for odd k, $L_k = \lceil (k+1)/2 \rceil (p+q)$, while for even k, $L_k = L_{k+1} - d$. We claim that the following lemma holds:

Lemma 10. *For a string $\boldsymbol{x}$ with $k \geq 0$ holes, if $\boldsymbol{x}$ is weakly p, q periodic and $|\boldsymbol{x}| \geq L_k$, then a substring of $\boldsymbol{x}$ of length at least L_0 is d periodic, where $d = \gcd(p, q)$.*

Proof. We prove this result by induction. For $k = 0$ and $k = 1$, the lemma holds by the periodicity lemmas for zero hole and one hole. If it holds for $k-2$, then for a string $\boldsymbol{x}$ with $|x| \geq L_k$, we consider its prefix $\boldsymbol{x_1} = x[1..L_{i-2}]$ and its suffix $\boldsymbol{x_2} = x[n - L_1 + 1..n]$ of length L_1. If the number of holes in $\boldsymbol{x_1}$ is less than or equal to $k - 2$, then by the inductive assumption $\boldsymbol{x_1}$ has a d periodic substring of length L_0. Otherwise the number of holes in $\boldsymbol{x_1}$ is greater than $k - 2$, so that the number of holes in $\boldsymbol{x_2}$ is at most 1, implying by Lemma 1 that $\boldsymbol{x_2}$ is d periodic. □

Note that unlike the 2-hole and 3-hole cases, in a string $\boldsymbol{x}$ with more than three holes the substring of $\boldsymbol{x}$ (let's call it $\boldsymbol{x}_d$) that may initially be d periodic is not necessarily a prefix or a suffix of $\boldsymbol{x}$. Therefore if $\boldsymbol{x}_d$ can be extended both to the left and to the right until all of $\boldsymbol{x}$ is covered, we may still claim that all of $\boldsymbol{x}$ is d periodic. Observe that $\boldsymbol{x}_d$ must itself contain a substring of length d without holes:

- in the case that $|\boldsymbol{x}_d| = L_0$, $\boldsymbol{x}_d$ contains no holes and $L_0 \geq 2d$;
- if $|\boldsymbol{x}_d| = L_1$, $\boldsymbol{x}_d$ contains at most one hole and $L_1 \geq 3d$.

Figure 3 demonstrates a possible position of x_d and a substring of x_d without holes.

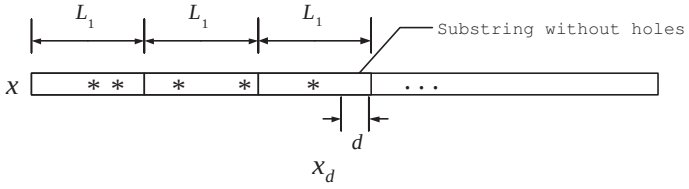

Figure 3: Possible position of x_d

To accommodate three or more holes, we give a more general definition of RE and LE as follows:

Definition 11. Suppose a string $\boldsymbol{x}$ with zero or more holes is weakly p, q periodic, $q > p$, with a substring $\boldsymbol{x}_d = \boldsymbol{x}[i..j]$, $j - i \geq p - 1$, that is d periodic, $d = \gcd(p, q)$.

(a) $\boldsymbol{x}_d$ is said to be RE iff $\boldsymbol{x}[j + 1] = \{\Sigma\}$ (hole) or there exists an integer sequence $s_1, s_2, \ldots, s_t$, $t \geq 2$, such that

- $s_1 = j + 1 \leq n$ and $s_t \in i..j$;
- for every $\ell \in 2..t$, $\boldsymbol{x}[s_\ell] \in \Sigma$ and $|s_\ell - s_{\ell-1}| = p$ or q.

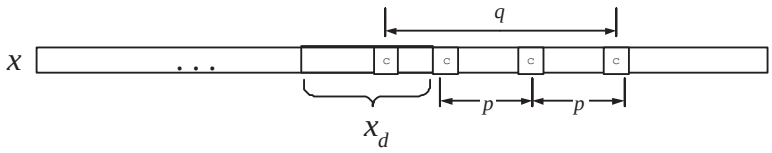

Figure 4: Example of RE and a path

(b) Symmetrically, $\boldsymbol{x}_d$ is LE iff $\boldsymbol{x}[i-1] = \{\Sigma\}$ or there exists an integer sequence $s_1, s_2, \ldots, s_t$, $t \geq 2$, such that

- $s_1 = i - 1 \geq 1$ and $s_t \in i..j$;
- for every $\ell \in 2..t$, $\boldsymbol{x}[s_\ell] \in \Sigma$ and $|s_\ell - s_{\ell-1}| = p$ or q.

Intuitively, this definition means that if we can find a path starting from $\boldsymbol{x}[j+1]$ that at each step identifies a next position p or q positions away and not a hole, terminating at a position that lies between i and j — then $\boldsymbol{x}[i..j]$ is RE (similarly for LE). Figure 4 illustrates an example of RE and such a path.

Note that Definitions 4 and 6 given in the previous section are special cases of this general definition.

Lemma 12. *Suppose that a string $\boldsymbol{x}$ with zero or more holes is weakly p, q periodic, $q > p$, with $d = \gcd(p, q)$. If there exist i and $j \geq i + p - 1$ such that $\boldsymbol{x}[i..j]$ is d periodic and RE (respectively, LE), then $\boldsymbol{x}[i..j+1]$ (respectively, $\boldsymbol{x}[i-1..j]$) is d periodic.*

Proof. We prove the RE case only. If $\boldsymbol{x}[j+1] = \{\Sigma\}$ then certainly for every $\ell \in i..j$ such that $\ell \equiv (j+1) \bmod d$, $\boldsymbol{x}[\ell] \approx \boldsymbol{x}[j+1]$. Otherwise there exists a sequence $s_1, s_2, \ldots, s_t$ as described in Definition 11(a). We see that

$$\boldsymbol{x}[j+1] \approx \boldsymbol{x}[s_2] \approx \boldsymbol{x}[s_3] \approx \cdots \approx \boldsymbol{x}[s_t],$$

and since every $\boldsymbol{x}[s_\ell] \in \Sigma$, $2 \leq \ell \leq t$, it follows that $\boldsymbol{x}[j+1] \approx \boldsymbol{x}[s_t]$. Since moreover $j + 1 \equiv s_\ell \bmod d$ for every $\ell \in 2..t$, we conclude in particular that $j+1 \equiv s_t \bmod d$. Since $s_t \in i..j$ and $\boldsymbol{x}[i..j]$ is d periodic, therefore $\boldsymbol{x}[j+1] \approx \boldsymbol{x}[r]$ for every $r \in i..j$ such that $r \equiv (j+1) \bmod d$. Thus $\boldsymbol{x}[i..j+1]$ is d periodic, as required. □

We now define functions ***Right-Extend*** and ***Left-Extend*** as follows:

Definition 13. Suppose that $\boldsymbol{x}$ is weakly p, q periodic, $q > p$, with a d periodic substring $\boldsymbol{x}[i..j]$, where $d = \gcd(p, q)$ and $j - i \geq p - 1$. The function ***Right-Extend*** maps the pair (i, j) to $(i, j+1)$ if $\boldsymbol{x}[i..j]$ is RE and to (i, j) otherwise. The function ***Left-Extend*** maps the pair (i, j) to $(i-1, j)$ if $\boldsymbol{x}[i..j]$ is LE and to (i, j) otherwise.

Using these functions, we can state a general characterization of the left and right extensions that guarantee that $\boldsymbol{x}$ is d periodic.

Lemma 14. *If $\boldsymbol{x}$ with $k \geq 0$ holes has weak periods p and $q > p$, and $|\boldsymbol{x}| \geq L_k$, then at least a substring $\boldsymbol{x}[i..j]$ of length L_0 is d periodic, where $d = \gcd(p, q)$. If there exists a concatenation of functions $E = E_1 \circ E_2 \circ \cdots \circ E_t$ where for every $\ell \in 1..t$, $E_\ell \in \{\textbf{\textit{Right-Extend}}, \textbf{\textit{Left-Extend}}\}$, and such that $E(i, j) = (1..n)$, then $\boldsymbol{x}$ is d periodic.* □

This is a statement of the periodicity lemma that applies to all strings with or without holes. However, as in the two-hole case (Corollary 9), we can be more precise: we now describe a straightforward algorithm that identifies a maximum-length d periodic substring of $\boldsymbol{x}$ that contains a substring intially known to be d periodic. The algorithm uses a list of the k holes in $\boldsymbol{x}$ and executes in $O(k)$ time.

Consider $\boldsymbol{x} = \boldsymbol{x}[1..n]$, $n \geq L_k$, with $k \geq 0$ holes. Suppose an array $H[1..k]$ gives the locations of all the holes in $\boldsymbol{x}$ in ascending order. We add $H[0] = 0$ and $H[k+1] = n+1$. By Lemma 10 we may suppose that a $\Theta(k)$ scan of H has yielded a range $i..j$ in $\boldsymbol{x}$ such that $\boldsymbol{x}[i..j]$ is d periodic, as well as a position s in H such that $H[s] < j$, $H[s+1] > j$, where in addition one of the following holds:

* $j-i > L_0$ and $H[s] < i$;
* $j-i > L_1$ and $H[s-1] < i$, $H[s] \in i..j$.

In either of these cases $\boldsymbol{x}[i..j]$ contains a substring $\boldsymbol{x}[\ell..\ell+d-1]$ such that for every $i' \in \ell..\ell+d-1$, $\boldsymbol{x}[i'] \in \Sigma$ (i' not a hole).

In addition to H, it is convenient also to compute a Boolean array $N[1..k]$ defined as follows: for every $s \in 1..k$, $N[s] =$ `TRUE` if $\boldsymbol{x}\big[H[s]+q-p\big]$ is a hole, $N[s] =$ `FALSE` otherwise. Figure 5 describes the preprocessing that computes N in $\Theta(k)$ time.

```
N[k] ← FALSE; r ← 2; δ ← q−p
for s ← 1 to k−1 do
    START ← H[s]
    while r ≤ k and H[r]−START < δ do
        r ← r+1
    if r > k or H[r]−START ≠ δ or H[r]+p > n then
        N[s] ← FALSE
    else
        N[s] ← TRUE; r ← r+1
```

Figure 5: Preprocessing: compute $N = N[1..k]$ in $\Theta(k)$ time from the array H of holes.

We are now in a position to describe an algorithm that extends a d periodic range $i..j$ in $\boldsymbol{x}$ to the right by processing H and N from left to right, with minimal access to $\boldsymbol{x}$ itself. The function `right-extend` shown in Figure 6 uses a current hole s to extend the current range: it returns $s+1$ and an extended right boundary j if further extension to the right may be possible; otherwise, it returns $s = k+1$ and the absolute rightmost boundary j of the

d periodic substring. It executes in constant time for each position s in H. (Note that here we assume the mathematical `mod` operation can be performed in constant time, since (a `mod` $b = a - \lfloor a/b \rfloor \cdot b$); thus the complexity of `mod` is equivalent to that of division and multiplication.) A corresponding algorithm `left-extend` deals with left extension of range $i..j$. Overall, repeated execution of `right-extend` and `left-extend` will yield a maximum-length d periodic substring that contains the original d periodic range $i..j$, thus generalizing the algorithm described in Figure 1 for the two-hole case.

function `right-extend`$(H, N, s, k, \boldsymbol{x}, i, j, \ell, n, p, q, d)$
if $j - H[s] \geq q$ **or not** $N[s]$ **then**
    $s \leftarrow s+1;\ j \leftarrow \max\{j, \min\{H[s]+q-1, n\}\}$
else
    $j \leftarrow H[s]+q$
    if $\boldsymbol{x}[j] \approx \boldsymbol{x}[\ell+(j-i) \bmod d]$ **then**
        $s \leftarrow s+1$
    else
        $j \leftarrow j-1;\ s \leftarrow k+1$
return j, s

Figure 6: This function uses a single hole $H[s]$ to extend the d periodic range $i..j$ to the right.

We remark that a little further preprocessing may be done to form an array $\boldsymbol{z}[1..d] = \boldsymbol{x}[\ell..\ell+d-1]$. Apart from H, N and $\boldsymbol{z}$, at most one reference to $\boldsymbol{x}[j]$ is then required in order to right-extend range $i..j$.

For a string $\boldsymbol{x}$ with multiple holes and with weak periods $p = 4$ and $q = 6$, we illustrate the right extend process in Tables 6–8. Starting in Table 6, we first identify a periodic substring $\boldsymbol{x}[1..10]$ of length $p+q = 10$ with strong period $d = \gcd(4, 6) = 2$. As we already know, the existence of such a substring is guaranteed by Lemma 10. Let $[\ell..\ell+d-1]$ be $\boldsymbol{x}[1..2]$. Since the position of the first hole $H[s] = 9$, we immediately know that $\boldsymbol{x}[1..9+q-1]$ is d periodic. Because every position in $x[9+1..9+q-1]$ is `RE` according to Definition 11, Lemma 12 tells us that $x[1..9+q-1]$ is d periodic. Since $N[s] =$ `TRUE` indicates that both $x[15-p]$ and $x[15-q]$ are holes, we have to compare $x[15]$ with $x[2]$. Since they match, we right-extend j from position 10 to 15.

Next we consider $H[s] = 11$ in Table 7. Since $N[s] =$ `FALSE`, without any comparison we know that $x[1..11+q]$ is d periodic and therefore right-extend j to position 17.

Finally we consider $H[s] = 12$ in Table 8. Because $N[s] =$ `TRUE` and $x[12+q]$ does not match $x[2]$, the algorithm correctly returns the maximum d periodic range 1..17.

5. Summary and Future Work

The periodicity lemma is perhaps the fundamental result of stringology. In this paper we extend this result to strings with holes, an increasingly important

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
$\boldsymbol{x} =$	a	b	a	b	a	b	a	b	$*$	b	$*$	$*$	a	$*$	a	b	a	c	a	b
	i									j										

Table 6: $H[s] = 9, N[s] = TRUE, x[15] \approx x[1], j \leftarrow 15$

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
$\boldsymbol{x} =$	a	b	a	b	a	b	a	b	$*$	b	$*$	$*$	a	$*$	a	b	a	c	a	b
	i														j					

Table 7: $H[s] = 11$ $N[s] = FALSE, j \leftarrow 17$

algorithmic topic. Throughout this paper we have used elementary and simple methods independent of number theory. In the case that the number of holes is arbitrary, we have taken a quite different approach than the graph-theoretical one of [BS04]. Our Lemma 14 is very general, covering indeterminate strings whose holes are not necessarily don't-cares; it leads to the algorithm that identifies maximum-length d periodic substrings of $\boldsymbol{x}$. We would like to extend other important results in stringology to strings with holes (indeterminate strings).

References

[BB99] J. Berstel and L. Boasson. Partial words and a theorem of Fine and Wilf. *Theoret. Comput. Sci.*, 218:135–141, 1999.

[BS04] F. Blanchet-Sadri. Periodicity on partial words. *Computers and Mathematics with Applications*, 47:71–82, 2004.

[BSH02] F. Blanchet-Sadri and Robert A. Hegstrom. Partial words and a theorem of fine and wilf revisited. *Theor. Comput. Sci.*, 270(1-2):401–419, 2002.

[FP74] M. J. Fischer and M. S. Paterson. String matching and other products. In R.M.Karp, editor, *SIAM-AMS Proceedings*, number 7 in Complexity of Computation, pages 113–125, 1974.

[FW65] N. J. Fine and H. S. Wilf. Uniqueness theorems for periodic functions. In *Proc. Amer. Math. Soc.*, volume 16, pages 109–114, 1965.

[HS03] Jan Holub and W. F. Smyth. Algorithms on indeterminate strings. In Mirka Miller and Kunsoo Park, editors, *Proc. 14th Australasian Workshop on Combinatorial Algorithms*, pages 36–45, 2003.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
$\boldsymbol{x} =$	a	b	a	b	a	b	a	b	$*$	b	$*$	$*$	a	$*$	a	b	a	c	a	b
	i																j			

Table 8: $H[s] = 12$ $N[s] = TRUE, not(x[18] \approx x[2]),$ `return`

[HSW06] Jan Holub, W. F. Smyth, and Shu Wang. Hybrid pattern-matching algorithms on indeterminate strings. *London Algorithmics and Stringology 2006, J. Daykin, M. Mohamed and K. Steinhoefel (eds.), King's College London Series Texts in Algorithmics*, pages 115–133, 2006.

[HSW08] Jan Holub, W. F. Smyth, and Shu Wang. Fast pattern-matching on indeterminate strings. *J. Discrete Algorithms*, 6(1):37–50, 2008.

[IMM+03] C.S̃. Iliopoulos, Manal Mohamed, Laurent Mouchard, Katerina G. Perdikuri, W.F̃. Smyth, and Athanasios K. Tsakalidis. String regularities with don't cares. *Nordic J. Computing*, 10(1):40–51, 2003.

[Smy03] Bill Smyth. *Computing Patterns in Strings*. Pearson Addison Wesley, 2003.