F-SINGULARITIES AND FROBENIUS SPLITTING NOTES 8-26-2010

KARL SCHWEDE

1. Assumptions and notation

Throughout all rings will be Noetherian and excellent. The excellent assumption can in many cases be removed, but for simplicity we will keep it.

Often rings will be assumed to contain a field of characteristic $p > 0$. If R is a ring of characteristic $p > 0$, it possesses that absolute Frobenius map $F : R \to R$. This is the map defined by $F(r) = r^p$ it is a map of rings. It thus turns R into an R-module with a non-standard action. That is, $r.x = r^p x$. We denote this R-module by $F_* R$. Why? Well, if $X = \operatorname{Spec} R$, then $F : \mathcal{O}_X \to F_* \mathcal{O}_X$ is the structural map associated to Frobenius. There are other common notations as well.

(a) ${}^1 R$.
(b) $R^{1/p}$ if R is reduced.

You may notice the number 1 in front of the R, and wonder why it's there. The point is that you can iterate Frobenius $F^e = F \circ F \circ \cdots \circ F$ and have induced module structures on R, denoted by $F^e_* R \cong {}^e R \cong R^{1/p^e}$. It is useful to observe that F^e_* is an *exact functor*.

These different notations for the same thing have different advantages. $R^{1/p}$ is useful because it allows one to easily distinguish elements from R and $F_* R$. On the other hand, it can lead to confusing statements since if we view $I^{1/p} \subseteq R^{1/p}$ as the ideal of $R^{1/p}$ made up of pth roots of I, then $(I^{1/p})^p = (I^p)^{1/p} \neq I$ (the latter is an ideal of R, where the two former are ideals of $R^{1/p}$). $I^{1/p}$ also is not a decent notation for modules.

Definition 1.1. Given an ideal $(x_1, \ldots, x_n) = I \subseteq R$, we use $I^{[p^e]}$ to denote the ideal $(x_1^{p^e}, \ldots, x_n^{p^e})$.

It is easy to see that this definition is independent of the choice of generators of I since $I^{[p^e]}$ can also be identified with the $F^e_* R$-ideal $I \cdot (F^e_* R)$.

Example 1.2. Consider the ring $R = \mathbb{F}_p[x_1, \ldots, x_n]$. Then $F_* R$ is a free R-module with basis $\{x_1^{\lambda_1} x_2^{\lambda_2} \ldots x_n^{\lambda_n} | 0 \leq \lambda_i \leq p - 1\}$.

The object $F_* R$ plays well with localization and completion.

Lemma 1.3. *Suppose that R is a ring of characteristic $p > 0$, $\mathfrak{m}$ is a maximal ideal and W is a multiplicative set. Then*

(i) $W^{-1}(F_* R) \cong F_*(W^{-1} R)$
(ii) $F_* \hat{R} \cong \widehat{F_* R}$ *(where the second is completion as an R-module).*

where the $\hat{\_}$ denotes completion with respect to $\mathfrak{m}$.

Proof. The first statement follows since $W^{-1}(F_* R) = F_*((W^p)^{-1}R)$ but $(W^p)^{-1}R \cong W^{-1}R$ since $r/w = (rw^{p-1})/w^p$. For (ii), notice first that $\hat{R} = \lim_{\leftarrow} R/\mathfrak{m}^n = \lim_{\leftarrow} R/(\mathfrak{m}^n)^{[p^e]}$ since the two sequences of ideals are cofinal. Then

$$\widehat{F^e_* R} = \lim_{\leftarrow}(F^e_* R)/\mathfrak{m}^n = \lim_{\leftarrow} F^e_*(R/(\mathfrak{m}^n)^{[p^e]}) = F^e_* \lim_{\leftarrow}(R/(\mathfrak{m}^n)^{[p^e]}) = F^e_* \lim_{\leftarrow} R/\mathfrak{m}^n = F^e_* \hat{R}.$$

□

Of course, there is another functor also, F^* which is defined by $F^*\mathscr{L} = \mathscr{L} \otimes F_*\mathcal{O}_X$ (and then viewed as an $F_*\mathcal{O}_X = \mathcal{O}_X$ module via the action on the right). Unlike F_*, F^* is not exact in general (although it sometimes is, as we will see). If $\mathscr{L}$ is a line bundle, then $F^*\mathscr{L} = \mathscr{L}^p$. One can see this by looking at the transition functions and noticing that they are raised to powers.

Definition 1.4. A ring of characteristic $p > 0$ is said to be *F-finite* if the Frobenius map is a finite map. In other words, if R is reduced, this means that $R^{1/p}$ is a finite R-module.

Lemma 1.5. *If R is F-finite, so is any quotient, localization, or completion at a maximal ideal.*

Proof. Suppose that R is F-finite, thus we have a surjective map of R-modules $\oplus_{i=1}^n R \to F_* R$ for some n. If W is a mutiplicative set then tensoring with $W^{-1}R$ will give us a new surjection. Completion is similar and quotienting out by an ideal is also straightforward. □

Note that thus if you start with a variety over an algebraically closed (or perfect) field, anything you might ever end up working with is still F-finite (even if you eventually move beyond having perfect residue fields) because $k[x_1, \dots, x_n]$ is F-finite (as long as k is an F-finite, eg perfect, field). The usual examples of non-perfect fields, $\mathbb{F}_p(x)$ are still F-finite! Although $\mathbb{F}_p(x_1, \dots, x_n, \dots)$ is not F-finite.

Technical lemmas we won't prove.

Lemma 1.6. [**?**][**?**] *If R is F-finite then R is excellent and it has a dualizing complex.*

Remark 1.7. If you don't know what a dualizing complex is, don't worry about it.

In other words, if you assume F-finite, you're working in a pretty geometric setting already.

2. Flatness of Frobenius

Suppose that R is a noetherian ring of characteristic $p > 0$. In [**?**], Kunz noticed the following: If $F_* R$ is flat as an R-module and $R \subseteq S$ is unramified in codimension 1, then $R \subseteq S$ is unramified.

Definition 2.1. An extension $R \subseteq S$ is called *unramified* is for every $\mathfrak{q} \in S$ with $\mathfrak{p} = \mathfrak{q} \cap R$, one has that $\mathfrak{p}S = \mathfrak{q}S$ and also that $k(\mathfrak{p}) \subseteq k(\mathfrak{q})$ is separable.

He then noticed that the condition that $F_* R$ is a flat R-module is equivalent to R being regular.

Theorem 2.2. [**?**] *Suppose that R is a local ring of characteristic $p > 0$. Then R is regular if and only if $F_* R$ is flat as an R-module.*

Proof. [**?**] We'll only prove the ($\Rightarrow$) direction today. We do not assume that R is F-finite. Suppose that R is regular, then $\hat{R}$ is a power series ring $k[[x_1, \ldots, x_n]]$ where k is the residue field of R. We have the following diagram:

$$\begin{array}{ccccc}
R & \xrightarrow{\beta} & \hat{R} & = & k[[x_1, \ldots, x_n]] \\
\Big\downarrow{\scriptstyle \alpha} & & & & \Big\downarrow \\
 & & & & k^{1/p}[[x_1, \ldots, x_n]] \\
 & & & & \Big\downarrow \\
R^{1/p} & \xrightarrow{\gamma} & \hat{R}^{1/p} & = & k^{1/p}[[x_1^{1/p}, \ldots, x_n^{1/p}]]
\end{array}$$

Once we show that the right vertical column is flat, then we know that $\gamma \circ \alpha$ is also flat. This combined with the fact that γ is faithfully flat implies that α is flat by [**?**, Page 46].

So, we need to show that the right vertical column is flat. The inclusion $k^{1/p}[[x_1, \ldots, x_n]] \subseteq k^{1/p}[[x_1^{1/p}, \ldots, x_n^{1/p}]]$ is clearly flat since the target is free as an R-module. The other inclusion is also free since $k^{1/p}$ is a flat k-module (this requires a little bit of work, Kunz cites [**?**, Chapter III, Section 5]). $\square$

Thus on a regular variety X, $F^e_* \mathcal{O}_X$ is a locally free sheaf (of F-finite rank assuming that X is F-finite). In particular, F^* is an exact functor if and only if X is regular.

Proposition 2.3. *If $X = \operatorname{Spec} R$ is an F-finite regular affine scheme, then $F^e : \mathcal{O}_X \to F^e_* \mathcal{O}_X$ splits as a map of $\mathcal{O}_X$-modules.*

Proof. First we claim that the statement is local. Indeed, consider the map $\sigma : \operatorname{Hom}_R(F^e_* R, R) \to R$ defined by evaluation at 1. The map F^e defined in the statement of the proposition splits if and only if σ surjects. The surjectivity of σ is a local property (since R is F-finite), so we can assume that R is local. Thus $F^e_* R$ is a flat and thus free R-module. Therefore there exist many surjective maps $\phi : F^e_* R \to R$ (project onto one component) we just need to see that one of them is a splitting. Suppose $\phi(x) = 1$ for some $\phi \in \operatorname{Hom}_R(F^e_* R, R)$ and some $x \in F^e_* R$, but then $\psi(\_\_) = \phi(x \cdot \_\_)$ clearly is a splitting of F^e. $\square$

The splitting of Frobenius is a statement about the singularities of X. If it occurs, it says something about the singularities being mild (we'll see some very effective criteria for checking this in a couple days).

Example 2.4. Let us compute $F^e_* \mathcal{O}_X$ on $X = \mathbb{P}^1_k$, where $k = \bar{k}$. We know that $F^e_* \mathcal{O}_X = \mathcal{O}_X(a_1) \oplus \mathcal{O}_X(a_2) \oplus \cdots \oplus \mathcal{O}_X(a_{p^e})$ because we are working on $\mathbb{P}^1$. We also know that $H^0(X, F^e_* \mathcal{O}_X) = k$ so exactly one of the a_i is equal to zero (and the rest are negative), say $a_1 = 0$. We will show that the rest of the $a_i = -1$, to see this consider

$$\begin{aligned} k^{p^e+1} &= H^0(X, (F^e_* \mathcal{O}_X(p^e))) = H^0(X, (F^e_* \mathcal{O}_X) \otimes \mathcal{O}_X(1)) \\ &= H^0(X, \mathcal{O}_X(a_1 + 1) \oplus \mathcal{O}_X(a_2 + 1) \oplus \cdots \oplus \mathcal{O}_X(a_{p^e} + 1)) \geq k^{2+(a_2+2)+\cdots+(a_{p^e}+2)}. \end{aligned}$$

But the only way this will happen is if each $a_i = -1$ for $i \geq 2$ (since they all already negative numbers).

For $X = \mathbb{P}^1$, we saw that $\mathcal{O}_X \to F^e_* \mathcal{O}_X$ is also going to split (because 1 goes to 1). However, not all smooth varieties which have locally split Frobenius have globally split

Frobenius. Projective space does (as we'll see, as do toric varieties in general and Fano varieties in "most" characteristics).

Example 2.5. Suppose that X is a supersingular elliptic curve, see [**?**, Chapter IV, Section 4, page 332], in other words $F : H^1(X, \mathcal{O}_X) \to H^1(X, F_* \mathcal{O}_X)$ is the zero map. Then X is not Frobenius split. To prove it, observe that $H^1(X, \_\_)$ is a functor. On the other hand, one can show that if X is an ordinary elliptic curve, it is Frobenius split (more on this later).

Frobenius split varieties satisfy strong properties.

Lemma 2.6. *Suppose that X is a variety whose Frobenius morphism splits. Then for any ample line bundle $\mathscr{L}$ on X, $H^i(X, \mathscr{L}) = 0$ for all $i \geq 0$.*

Proof. Note that $\mathcal{O}_X \to F^e_* \mathcal{O}_X$ splitting implies that $\mathscr{L} \otimes \mathcal{O}_X \to \mathscr{L} \otimes F^e_* \mathcal{O}_X = F^e_*(\mathcal{O}_X \otimes (F^e)^* \mathscr{L}) = F^e_* \mathscr{L}^{p^e}$ also splits. We then have $H^i(X, \mathscr{L}) \to H^i(X, F^e_* \mathscr{L}^{p^e})$ injects. But the right side vanishes by Serre vanishing for $e \gg 0$ so thus the left side vanishes too. □

Even though Kodaira vanishing fails in positive characteristic, it holds for Frobenius split varieties.

Theorem 2.7. *Suppose that X is a projective Frobenius split variety. Then for any ample line bundle $\mathscr{L}$ on X, $H^i(X, \omega_X \otimes \mathscr{L}) = 0$ for $i > 0$.*

Proof. It's not hard, but we'll prove it a little later. □

We also briefly mention a link to projective normality.

Definition 2.8. Suppose that $Y \subseteq X$ is a closed subvariety of X. Given a map $\phi : F^e_* \mathcal{O}_X \to \mathcal{O}_X$, we say that Y is *ϕ-compatible* if ϕ induces a map $\bar{\phi} : F^e_* \mathcal{O}_Y \to \mathcal{O}_Y$ by restriction.

Theorem 2.9. *If $\phi : \mathcal{O}_{\mathbb{P}^n} \to \mathcal{O}_{\mathbb{P}^n}$ is a splitting of Frobenius, then any ϕ-compatible normal $Y \subseteq \mathbb{P}^n$ is embedded in $\mathbb{P}^n$ projectively normally.*

Proof. It is sufficient to show that $H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(i)) \to H^0(Y, \mathcal{O}_Y(i))$ is surjective for all i (see [**?**, Chapter II, Exercise 5.14]). Consider the following commutative diagram:

$$\begin{array}{ccc} H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(p^e i)) & \xrightarrow{\phi(i)} & H^0(\mathbb{P}^n, \mathcal{O}_{\mathbb{P}^n}(i)) \\ \gamma \downarrow & & \downarrow \delta \\ H^0(\mathbb{P}^n, \mathcal{O}_Y(p^e i)) & \xrightarrow{\bar{\phi}(i)} & H^0(\mathbb{P}^n, \mathcal{O}_Y(i)) \end{array}$$

By Serre vanishing, γ is surjective and $\bar{\phi}(i)$ is also surjective because it is induced from a splitting. Thus δ is surjective as well. □