Responses to Reviewer 2

Comment 1:

Wang et al. (2023, Fire) already compared coupled and uncoupled simulations of a fire in the grassland-forest transition zone near the Mongolia-Inner Mongolia border, using the same atmospheric grid spacings and FROM-GLC-derived fuels. What does the present study add? I feel this needs a clearer answer in the Introduction. The manuscript already introduces the FireFlux and FireFlux II evaluations, but why are the present coupling results never compared with those studies? Section 3 contains no literature citations. Rothermel (1972) and Coen et al. (2013) are also absent, despite their relevance to the model description.

Response to comment 1:

The coupled and uncoupled simulations corroborate Wang et al. (2023). In that study, two-way coupling increased the burned area by about 17% after roughly 130 h. Here, the coupled burned area likewise grows faster after about 12 h. The present runs also show a later stage: after about 17 h the feedback has modified the local wind and thermal structure, further expansion in the coupled simulation is limited, and the uncoupled burned area can become larger. **The scope of the study is also different. It compares meteorological forcings and land-cover-derived fuel maps, and it quantifies the responses to meteorological and fuel perturbations.**

The coupling results were not compared with FireFlux and FireFlux II because the observations are of a different kind. Those were instrumented experimental burns of a few hectares, lasting minutes, with tower measurements of wind, temperature and rate of spread. The Chenbarhu fire is evaluated with the final burn scar and the surveyed area. There is no tower record and no time-resolved perimeter, so the coupled-uncoupled difference cannot be scored on the same quantities.

Section 3 reports the experimental results, and the submitted text therefore did not cite the literature those results relate to. We will add the citations where the results are interpreted. **Rothermel (1972) and Coen et al. (2013) will be added in Section 2.2**, for the spread equation and the two-way coupling, and cited again in Section 3.2 for the responses to wind, fuel moisture and fuel load. The coupled-uncoupled comparison in Section 3.1 will cite Wang et al. (2023).

Comment 2:

I could not find complete namelist.input in the materials provided. L163-165 state that the fire-driving wind is interpolated to 6.096 m. Does this match the actual fire_wind_height setting, and why was this height chosen for the grassland fuels? Was an additional wind-reduction factor applied? Please also identify the fire module, spread formulation and effective upwinding scheme. Given the sharp fuel transitions and non-burnable patches, were numerical propagation errors near fuel boundaries checked?

Response to comment 2:

We thank the reviewer for this careful question. It concerns the wind that actually enters the spread calculation, and the submitted manuscript did not state that configuration

clearly enough. The complete **namelist.input** used for the simulations will be added to the supplement. That file does not set fire_wind_height, fire_wind_log_interp, or fire_use_windrf, so the code defaults are used. These defaults will be stated in Section 2.2. The height of 6.096 m given at L163-165 is the height used in the runs. fire_wind_height defaults to 6.096 m. With fire_wind_log_interp = 1, the wind on a burnable cell is log-interpolated to the height stored for its fuel category, and that height is 6.096 m for all 13 Anderson fuel models, including short grass. This is the 20 ft reference height in the Rothermel fuel table. It was retained from the code default and was not chosen separately for the grassland. No additional wind-reduction factor was applied. fire_use_windrf defaults to 0, so the interpolated wind is not multiplied by the tabulated factor windrf. For short grass that factor is 0.36. The wind at 6.096 m therefore enters the Rothermel spread calculation directly, and the absolute spread rates correspond to this wind. The same choice is used in every experiment. Differences among the meteorological forcings, the coupling options, the land-cover maps, and the perturbations in Tables 4 and 5 are differences under this common setting. On the innermost domain the fire module is SFIRE (ifire = 2). The spread rate is given by the Rothermel formulation (ibeh = 1). The level-set gradient is evaluated with the ENO scheme (fire_upwinding = 3), with fire_viscosity = 0.4.

We did not make a separate test of numerical propagation across fuel boundaries. Non-burnable cells are fuel category 14. For that category the no-wind spread coefficient and the wind and slope coefficients are set to zero, so the formulated rate of spread is zero. The distance by which the level-set front may advance numerically into those cells was not quantified.

Comment 3:

What is the source of the ignition time and coordinates at L124-126? Were 07:30 UTC and 119.09°E, 49.20°N used directly for model initialization? I would also like to know the ignition geometry, radius and duration.

Response to comment 3:

The ignition time and location were provided by the Hulun Buir Meteorological Bureau. The time is 07:30 UTC on 1 May 2023. The coordinates printed in the text, approximately 119.09°E and 49.20°N, are that location rounded to 0.01°. These two values were not entered as a single point, and the simulation was not initialized at 07:30 UTC. **The atmospheric integration starts at 00:00 UTC on 1 May 2023.** The fire is introduced only on the innermost domain.

The ignition is one straight line, with endpoints at 119.0861°E, 49.2042°N and 119.0871°E, 49.2052°N. The line is about 133 m long, and the ignition radius is 35 m, so a fire-grid cell within 35 m of the line can be ignited. This radius is larger than the 30 m fire-mesh spacing. The ignition start and end are both 27000 s after 00:00 UTC, which is 07:30:00 UTC. Because the two times are equal, the whole line is introduced at that instant rather than being drawn from one end to the other. The ignition rate of spread was left at its default of 0.01 m s^{-1}, so the ignited region grows outward from the line until it reaches the 35 m radius, which takes 3500 s. Spread after that time is

calculated with the Rothermel formulation.

Comment 4:

How exactly was wind speed perturbed in Table 4? Which vertical levels, nested domains and input files were changed? Please show how much the wind driving the fire actually changed. Can the wind-slope explanation at L405-407 also be quantified?

Response to comment 4:

The reviewer is right to separate the wind that was altered from the wind that reached the fire. The manuscript did not make that distinction, so the ±2 and ±4 m s^{-1} in Table 4 can be read as a direct change in the fire-driving wind. The change was made at the start of the calculation, in the environmental wind supplied to WRF, rather than in the wind used by the spread formula. The fire-driving wind is obtained only after WRF has integrated that environmental wind over the terrain. Adjusting it directly would detach it from the flow that produced it, including the turning and acceleration where the terrain blocks the wind, and the size of the adjustment would then be difficult to interpret.

The perturbations were applied to the WPS intermediate files met_em.d01, met_em.d02, and met_em.d03 for every forcing time. In each file, the horizontal wind components UU and VV were changed on all 34 metgrid levels. No other variable was changed, and the WRF input and boundary files were then generated from these modified met_em files. The wind speed was increased or decreased by 2 or 4 m s^{-1} while the wind direction was kept. Where the original speed was smaller than a decrease of 2 or 4 m s^{-1}, this calculation reverses the wind vector. On the innermost domain at 06:00 UTC, the lowest metgrid level did not contain such points, and its domain-mean speed changed from 10.49 m s^{-1} by exactly +2, −2, +4, and −4 m s^{-1}. That increment is not the change in the wind used for fire spread. In the area 119.0-119.2°E, 49.1-49.4°N, the mean fire-grid wind at 08:00 UTC was 6.21 m s^{-1} in the baseline and 6.28, 7.34, 6.82, and 7.15 m s^{-1} for the +2, −2, +4, and −4 m s^{-1} experiments. At 13:00 UTC the same means were 9.04 m s^{-1} and 8.14, 9.19, 6.22, and 8.50 m s^{-1}. The 10 m wind in this area changed in the same way.

The sentence at L405-407 refers to the same pathway. The perturbed environmental wind flows over the fire area, where the slopes are about 3°-5°. Where the terrain blocks the flow, the near-surface wind can turn or accelerate before it reaches the fire head, and the head moves with that wind. The wind change is the fire-grid change given above. No flat-terrain experiment was added, so the part caused only by terrain blocking is not given a separate percentage. Figure 8a shows the combined result: with a stronger imposed wind the head shifts toward the northeast, and with a weaker imposed wind the main axis turns toward the northwest and extends westward.

Comment 5:

What fuel moisture content was prescribed? Was the default 0.08 used, or were different values assigned? AC1 explains that RH affects spread indirectly through changes in the simulated atmosphere while fuel moisture remains fixed. I can follow that argument, but I do not see the diagnostics needed to explain the 85.4% range in Fig. 8c. How

much did the fire-grid winds change, and where? For the temperature perturbations, was specific humidity or RH held fixed?

Response to comment 5:

The fuel moisture content was the default value, 0.08, in every experiment in Table 4. The namelist does not replace fuelmc_g, fmoist_run remains false, and fire_fmc_read remains 1, so this single prescribed value is used on the whole fire grid. The fuel-moisture model is used only in Table 5. The relative-humidity experiments therefore do not change the moisture term in the spread formula. The 85.4% range in Fig. 8c is the response of the spread rate under that fixed moisture. The relative-humidity change was introduced in the environmental fields. RH was modified in met_em.d01, met_em.d02, and met_em.d03, on all 34 metgrid levels and at every forcing time, while the temperature in those files was left unchanged. At 06:00 UTC on the innermost domain, the lowest-level mean relative humidity changed from 18.1% by +10, −10, +5, and −5 percentage points. WRF then integrates this atmosphere, and the fire responds to the wind that results.

That wind changed in the direction expected for this pathway. In the area 119.0-119.2°E, 49.1-49.4°N, the mean fire-grid wind at 13:00 UTC was 9.04 m s^{-1} in the baseline and 7.90, 9.32, 8.41, and 9.42 m s^{-1} for the +10, −10, +5, and −5 percentage-point experiments. A moister atmosphere gave a weaker fire-grid wind, and a drier atmosphere gave a stronger one. The spread rate in Fig. 8c is the growth of the fire over the simulation, and the wind term in the Rothermel formulation depends nonlinearly on wind speed, so the percentage change in spread rate is larger than the percentage change in the area-mean wind. The 85.4% is that spread-rate response.

For the temperature perturbations, only TT was changed, and the RH field was held fixed. The same lowest-level mean relative humidity remained 18.1% when the temperature was shifted by ±2 and ±4 K. These met_em files contain relative humidity and temperature rather than specific humidity, and WRF diagnoses the water vapor from those two fields. The temperature experiments therefore remain temperature experiments: the imposed relative humidity is unchanged, and the weaker spread response to temperature already reported in Fig. 8 is the result obtained with the fuel moisture fixed at 0.08.

Comment 6:

Why is no meteorological evaluation presented? The paper quotes weather observations, but I could not find a comparison of the simulated wind, temperature or humidity with the Hailar station or other available measurements.

Response to comment 6:

We are grateful for the close reading of the fire-weather passage in Section 2.1. The manuscript already names Hailar and the conditions on 1 May 2023, and the baseline simulation belongs in that same passage, where the present text leaves it out. **We will supply it. The baseline fields were taken at the grid point nearest Hailar (49.25°N, 119.7°E), about 45 km east of the ignition.** The 10 m wind there is 9.9 m s^{-1} at 08:00 UTC and 8.4 m s^{-1} at 13:00 UTC, from the south-southeast, beside the station mean of

9 m s^{-1} cited in the manuscript. The 2 m temperature at that point is 20.5°C at 08:00 UTC and 19.7°C at 13:00 UTC, and the 2 m relative humidity is about 16% at both times, in the same dry range as the 12-13% given for the fire day. **In the revision, Section 2.1 will show wind speed, wind direction, 2 m temperature and humidity at this grid point through the simulation window, set against the Hailar record from the Hulun Buir Meteorological Bureau from which the cited 9 m s^{-1} mean was taken, with the station temperature from that same record placed next to the simulated 2 m temperature.** The 14.9-15.4 m s^{-1} wind and the 19 m s^{-1} gusts reported by Zhou et al. (2025) remain identified as the estimate for the fire area. With the station record and the simulation written side by side, the weather used to frame the fire experiments can be read together with the burn-scar comparison that follows.

Comment 7:

How are spread distance and flame length calculated, over which part of the fire, and at what time? The ROS curves in Figs. 4b and 7b look like spatial averages. Are they perimeter averages? "Variability" in Figs. 8c and 9b also needs a definition.

Response to comment 7:

Each point in Figs. 4b and 7b is the mean of ROS_FRONT on the fire front at that time. ROS_FRONT is written only on the 30 m fire-mesh cells the front is crossing, where FIRE_AREA is greater than 0 and less than 1. The burned interior and the fuel that has not yet ignited are excluded. Because the cells have equal area, the curve is the mean along the front, taken from the hourly output from ignition through the end of the simulation. Flame length uses that same front. On those cells SFIRE evaluates Byram's relation, $L = 0.0775\ (IR)^{0.46}$, where R is the local rate of spread and I is the fireline-intensity factor diagnosed from the reaction intensity, with L in metres. The flame length in Figs. 8 and 9 is the mean of the positive values on this front at 13:00 UTC. Spread distance is taken at the same time and is the greatest distance from the ignition to any burned cell, so it measures the farthest point the fire has reached. In the baseline these three values are a front-mean spread rate of 1.16 m s^{-1}, a mean flame length of 0.22 m, and a spread distance of 97 km. The variability in Figs. 8c and 9b is, for one metric and one factor, the difference between the maximum and the minimum across the perturbation members, divided by the baseline and expressed as a percentage. For fuel load, the flame-length increase of 27.3% at 1.75 times the baseline and the decrease of 54.5% at 0.25 times the baseline sum to the 81.8% range given in the text. The quantity is the spread among the perturbation members.

Comment 8:

AC1 states that 1299 km² comes from a field survey. What is the area of the digitized scar polygon, and which area enters the calculation of recall, precision and IoU? The manuscript mentions BAI and visual editing, but I would like more detail on how the two satellite images were used and how uncertain the resulting boundary is. AC1 also suggests that suppression explains the missing eastern part of the scar. What records or other evidence support that explanation?

Response to comment 8:

The scar area and the three scores come from the same polygon. The pre-fire image is the Level-1A scene from the GF-6 WFV on 26 April 2023, and the post-fire image is the Level-1A scene from the GF-1 WFV1 on 10 May 2023. Both were radiometrically calibrated and both have a resolution of 16 m. The two scenes were mosaicked into one coverage of the fire, extending from about 118.76°E to 119.58°E and from 49.20°N to 50.00°N. The burned extent was extracted from that mosaic with the Burned Area Index and the boundary was then edited visually, giving the final scar polygon. The area of this polygon is 1299 km², or 129900 ha, the same area reported by the post-fire survey of the local forestry and grassland authority. Recall, precision and IoU all use this polygon. Recall is the overlap area divided by 1299 km². For FNL 6 h (FD) the overlap at 13:00 UTC is 507.7 km², so the recall is 39.1%. Precision is the overlap area divided by the simulated burned area of that experiment. IoU is the overlap area divided by the sum of the simulated burned area and 1299 km², minus the overlap. The boundary is the visually edited Burned Area Index outline on the 16 m mosaic, and the 1299 km² is the area of that outline.

The eastern portion remains inside this final polygon, and none of the experiments burns it. The polygon is the extent after suppression. The model is integrated forward from the ignition and contains the meteorological fields, the fuel, the terrain and fire-atmosphere feedback, and it contains no firebreak and no control line. In Fig. 4 every experiment has a burned area larger than 1299 km² by 18:00 UTC. The eastern difference therefore corresponds to suppression that the model does not represent. No suppression coordinates are available to lay on the eastern boundary. The revision will state that the polygon is the final scar after suppression, and the missing eastern part will not be drawn as a located firebreak.

Comment 9:

Were VIIRS or Himawari-9 active-fire detections checked for this event? If usable detections exist, could they help constrain the fire progression or estimated arrival times? I also think the abstract should acknowledge that the simulations in Fig. 4 end with burned areas several times larger than the surveyed area, together with the limitation that suppression is omitted.

Response to comment 9:

It is helpful to have the end state of Fig. 4 stated in the abstract. The revised abstract will say that every simulation in Fig. 4 reaches a burned area several times the surveyed 1299 km² by 18:00 UTC, because the model is integrated forward from the ignition with the meteorological fields, the fuel, the terrain and fire-atmosphere feedback, and with suppression omitted. The spatial comparison with the scar remains the overlap at 13:00 UTC, during the stage of spread without suppression.

VIIRS and Himawari-9 active-fire detections were not used for this case. The fire extent is the final polygon extracted from the mosaicked GF-6 and GF-1 scenes and edited visually, with the surveyed area of 1299 km². An active-fire detection marks a pixel that was hot at one overpass. A series of such points can show which locations the fire had

reached, but it does not supply a continuous front and does not replace this final polygon. The scores in the paper stay with the polygon, and the added abstract sentence records that, without suppression, the simulated area grows beyond the surveyed area.

Comment 10:

I am not yet convinced by the explanation at L548-550 that feedback enhances spread by strengthening local winds and promoting low-pressure development. Where are those changes demonstrated? AC1 reiterates the explanation that feedback later limits expansion, allowing the noFD area to exceed the FD area, but this also needs evidence. FD minus noFD fields of wind at the fire wind height, pressure perturbation and heat flux around 12:00 UTC and after 17:00 UTC would help. Do the fire heads encounter the northern river or approach the d03 boundary around the crossover?

Response to comment 10:

The care taken to separate the earlier enhancement from the later crossing of the area curves is appreciated. Both stages can be read from the FD minus noFD fields in the FNL 6 h pair. At 12:00 UTC the burned area is 685 km² in FD and 517 km² in noFD. Near the head, the fire-grid wind that SFIRE uses for spread is 11.3 m s^{-1} in FD and 8.3 m s^{-1} in noFD. Both point north, and the wind at the head is stronger once the fire is coupled. The ground-fire heat flux returns to the atmosphere only when feedback is on: GRNHFX near the head is about 619 W m^{-2} in FD and zero in noFD. At the same place the surface pressure in FD is lower by 16 Pa on average, and 41% of the cells are lower by more than 50 Pa. The head is still 35 km south of the northern boundary of d03. A stronger local wind, a pressure deficit around the head, and heat flux returned to the atmosphere all occur while the FD area is growing faster. Lines 548-550 summarize this mechanism.

After 17:00 UTC the same feedback is still operating, and the area curves cross. The areas are 3546 km² and 3522 km² at 17:00 UTC, and 3911 km² and 4147 km² at 18:00 UTC. At 17:00 UTC the surface pressure near the head is still 36 Pa lower in FD, GRNHFX is still present only in FD, and the fire-grid wind is still stronger in FD, 9.9 m s^{-1} against 9.2 m s^{-1}, and 10.4 m s^{-1} against 7.6 m s^{-1} at 18:00 UTC. The feedback continues to modify the wind and the thermal structure at the head. The FD area no longer increases at the same rate, while the noFD fire keeps expanding under the background flow, so the noFD area can become larger. These are the earlier and later stages of the same coupling: spread is enhanced at first, and further outward expansion in FD is limited later. At the crossing, both heads are 1-3 km from the northern boundary of d03, while the eastern and western boundaries remain tens of kilometres away. Within about 1 km inside that northern boundary, non-burnable fuel covers about 3% of the cells and the rest is mostly short grass, so the head does not meet a river laid across the north. Both experiments face the same boundary. The difference between them remains whether feedback is on, together with the wind, pressure and heat flux that follow from it.

Comment 11:

I have several questions about the runtime tests. Table S2 changes the floating-point

model and target flags alongside compiler version, so I do not think the differences can be attributed to the compiler version alone. How variable were the three repeated timings? What domain decomposition was actually used in each hybrid and asynchronous I/O case?

Response to comment 11:

The three timings are clearer when read against Table S2. The three rows are the compilation systems that belong to the three compiler versions. v18 uses -O3 -xCORE-AVX512 -fp-model fast=2, with -qopenmp on the C flags and -heap-arrays 8192 on FCOPTIM. v21 uses -O3 -xHost -fp-model fast=2, together with -qopt-zmm-usage=high. v22 uses -O3 -xHost -fp-model precise, together with -qopenmp-simd. Changing the version brings the instruction-set target and the floating-point model that go with that version. The supplement states the same point: the optimization flags are set together with the compiler version. The same WRF-SFIRE source was built with each of these three systems. The mean times of three one-hour integrations are 17 min 39 s, 18 min 04 s and 17 min 59 s, so the latter two are 25 s and 20 s longer than v18. That difference is the difference among the three versions together with the compilation system of each version. Tables 7 to 10 stay with v18 and its compilation system. **Every tabulated time is the mean of three repetitions of the same one-hour case, and the three raw times will be placed beside the corresponding means.**
The two-dimensional layouts for the hybrid and asynchronous I/O tests are not printed in the supplement, and the namelist does not set nproc_x or nproc_y. WRF assigns the nearest-to-square layout of the compute ranks, with nproc_x no greater than nproc_y. A run with 320 compute ranks is recorded as 16×20, which follows that rule. Table 8 keeps 10 nodes and 400 cores. The four rows are 40 MPI processes per node with 1 thread, 400 ranks laid out 20×20; 20 processes with 2 threads, 200 ranks laid out 10×20; 10 processes with 4 threads, 100 ranks laid out 10×10; and 5 processes with 8 threads, 50 ranks laid out 5×10. Table 10 keeps 400 MPI processes in total. The number of I/O processes is the product of nio_tasks_per_group and nio_groups, equal to 0, 10, 20, 30, 40, 50, 60 and 70, and **the compute ranks are 400 minus that number**. The layouts are 20×20, 15×26, 19×20, 10×37, 18×20, 14×25, 17×20 and 15×22. The grouping tests keep the I/O processes at 50 and the compute ranks at 350, so the layout stays 14×25 and only the grouping of those 50 processes changes. The shortest time is 25 processes per group and 2 groups. **These values of nproc_x×nproc_y will be added to Tables 8 and 10.**

Comment 12:

In Table 9, the (15×8) configuration is about 17% slower than (8×15), although both have the same η. The paper recognizes a directional dependence, but why attribute it to the north-south orientation of the fire? Similar benefits from asymmetric decomposition and I/O changes have been reported for atmospheric WRF (Huang et al., 2023, Applied Sciences). A comparison with the fire module switched off would help test this explanation. Please also give the history interval and output volume, and confirm that the optimized configurations preserve the simulated fire and atmospheric results.

Response to comment 12:

The direction at a fixed asymmetry needs a separate statement. η measures only how far the two process counts differ. Both (8×15) and (15×8) have η = 0.88, and the runtimes are 22 min 48 s and 26 min 43 s, so the second is about 17% longer. On each of the three domains the number of grid points east-west equals the number south-north. In this case the fire extends northward from the ignition, and the long axis of the burned area and of the fire front lies south-north. The extra fire-module work and the fire-grid halo exchange lie along that south-north band. Assigning more processes south-north cuts this band into more segments, so the fire-grid work on each process is closer from one process to the next. Assigning more processes east-west leaves that south-north band on fewer processes. At the same η, the layout with more south-north processes is therefore faster. The (6×20) layout takes 22 min 02 s and the (20×6) layout takes 29 min 20 s, in the same sense. Huang et al. (2023) report that asymmetric decomposition and I/O changes can likewise shorten the runtime of atmospheric WRF. Table 9 is the directional difference measured with the fire module on and with the fire extending south-north, and it stands alongside that atmospheric result.

The history interval is 60 min for d01 and d02 and 10 min for d03, with one file at each output time. From 00:00 to 01:00 UTC this writes 11 wrfout files, 3.3 GB in total, of which d01 is 0.11 GB, d02 is 0.22 GB and d03 is 2.97 GB. The timing tests use this same one-hour case and the same history output. Tables 7 to 10 stay with v18 and its compilation system. The physics options, the initial fields and the fire module remain unchanged, and the changes are the process layout and the I/O path. These optimized configurations keep this fire and atmospheric simulation, and the differences in the tables are wall-clock times.

Comment 13:

Which fire-induced circulations is the 300 m grid expected to resolve, and why is Kain-Fritsch active on the 1.5 km domain?

Response to comment 13:

On the 300 m domain the boundary-layer scheme and Kain-Fritsch are switched off, and Table 1 uses three-dimensional diffusion with the 1.5-order turbulent-kinetic-energy closure, which is the large-eddy simulation on that domain. The circulations this grid is expected to resolve are the fire-induced motions on scales of about a kilometre and larger: the inflow beside the fire line, the plume-scale updraft, and the larger turbulent eddies within the plume. These circulations change the near-surface wind and the pressure around the fire line, and that change feeds back onto the spread. Flame-scale motions remain inside the fire module.

The 1.5 km domain cannot carry that large-eddy simulation. Its grid spacing does not resolve the plume circulations just described, and it is already fine enough to grow unstable grid-scale updrafts on its own. The YSU boundary-layer scheme and the Kain-Fritsch cumulus scheme are therefore kept on this domain, so that the convection it cannot resolve is parameterized and the parent domain stays stable. The 300 m nest then receives a stable driving state at its lateral boundary. On the 300 m domain both schemes

are off, and the fire-induced circulation is computed once by the large-eddy simulation rather than being computed again by the cumulus scheme.

Comment 14:

At L65-67, why describe fire-spread models as explicitly simulating chemical processes? Please qualify this statement for the semi-empirical spread formulation used here. The Introduction mentions chemistry in several places (e.g. L65-67 and L102-105), yet the study does not appear to have the chemistry module enabled. Why is that?

Response to comment 14:

The chemistry wording in the introduction is clearer when set beside the results this paper actually presents. Lines 65-67 refer to fire-spread models as a class of methods. The spread calculated here uses the semi-empirical Rothermel equation: the rate of spread is determined by the reaction intensity, the wind and slope terms and the fuel moisture, and combustion is represented by the sensible and latent heat returned to the atmosphere. That equation does not solve a chemical mechanism and does not track chemical species. The sentence will be revised to state that spread in this paper is computed from this semi-empirical equation, with combustion represented by the reaction intensity and the heat flux returned to the atmosphere.
Lines 102-105 describe the further coupling of WRF-SFIRE with WRF-Chem. That coupled simulation has been carried out, and it has produced emissions and particulate results. The question, figures and conclusions of this paper stop at two-way fire-atmosphere feedback, meteorological forcing, land cover, fuel parameters and the computational configuration, so those chemistry results are not shown here. The introduction will keep this coupling capability, and sentences that can be read as an analysis of the chemistry results in the present paper will be tightened, so that they identify the chemistry simulation as material for a later presentation rather than as content of the figures in this paper.

Comment 15:

I do not think the distinction between WRF-Fire and WRF-SFIRE at L163-165 is accurate as a general statement. Please check this passage.

Response to comment 15:

The sentence is clearer once it is set against the wind algorithm. Lines 163-165 correspond to the two vertical treatments in Mandel et al. (2011, Section 5.2). Both serve the wind that the Rothermel equation uses at 6.096 m, the 20 ft height of BEHAVE. The first treatment starts from the wind at a single reference height and multiplies by the fuel wind-reduction factor. For a surface wind that reference height is commonly 10 m, and the reduction then supplies the midflame wind. The WRF-Fire package carried in the WRF release follows this reference-height reduction. The default in WRF-SFIRE does not start from that single reference wind. It reads the horizontal wind on the model levels, the first of which is the lowest model level, and interpolates it to 6.096 m on the logarithmic profile (Mandel et al., 2011, Equation 23 and the algorithm that follows). The namelist in this study does not reset fire_wind_height, fire_wind_log_interp or

fire_use_windrf, so the default is used: fire_wind_log_interp = 1, the height given by the fuel category is 6.096 m, and the wind-reduction factor is not applied. Lines 163-165 will add Mandel et al. (2011, Section 5.2), so that the 10 m wind and the lowest model level are the two treatments, and will state that this study uses the second.

Comment 16:

Why is the burned-area response described as "consistent" at L411-413 when weaker winds and lower temperatures also increase burned area in Fig. 8b?

Response to comment 16:

There, "consistent" means that the three perturbations that favor burning point the same way: an increase in wind speed, a decrease in relative humidity and an increase in temperature each lead to a larger burned area. A weaker wind and a lower temperature can also enlarge the burned area, and that result stands together with the sentence. The sentence then gives the magnitude: an increased wind and a higher temperature have the greater effect on burned area. The increased-wind side lies above the decreased-wind side. For temperature, against the baseline of 1337.63 km², +4 K is 1649.80 km², while −2 K and −4 K are 1364.01 km² and 1375.90 km², so the increase at +4 K is much larger than the increase in the two cooler members. A larger burned area under a weaker wind and a lower temperature therefore stands together with the shared direction named by "consistent," and the further point of the original sentence is that the increased wind and the higher temperature have the greater effect.

Comment 17:

At L430-431, the fuel-moisture experiment changes ROS by -20.0% and flame length by -4.5%. Why is the response described as mainly affecting intensity-related metrics?

Response to comment 17:

Writing the four changes together in one sentence has made the passage difficult to read, and that difficulty is regretted. The wording "mainly affects intensity" at lines 430-431 follows from dividing the four changes into two groups. The spatial group comprises burned area and spread distance, which change by −4.7% and +2.5%, both small. The intensity group comprises spread rate and flame length, which change by −20.0% and −4.5%. Fireline intensity rises and falls with the spread rate, and flame length is obtained from that intensity, so the spread rate and the flame length belong to the intensity group. The fuel-moisture model mainly changes the intensity group, while the change in spatial extent remains limited. Lines 430-431 will name the metrics in each group and will keep this comparison.

Comment 18:

The recommendation of a 90-95% per-node load goes beyond the reported tests, since no 95% case is presented. I also could not find memory measurements supporting the claim of reduced memory pressure at L567-569.

Response to comment 18:

Two rows were removed when Table 7 was prepared for the preprint, so the 95% case does not appear there, and the difficulty this caused is regretted. The core-density test was run on 10 nodes with 40 cores per node. The cases actually run used 40, 38, 36, 34, 32 and 30 processes per node, corresponding to loads of 100%, 95%, 90%, 85%, 80% and 75%, with runtimes of 17 min 39 s, 17 min 13 s, 17 min 11 s, 17 min 16 s, 18 min 12 s and 19 min 12 s. The 95% case is the 38-process row and the 90% case is the 36-process row, and these two have the shortest runtimes. The range 90%-95% is the result of these two tests. The 38-process and 34-process rows will be restored to Table 7.
The memory-pressure statement refers to the hybrid runs in Table 8. Pure MPI places 40 processes on each node, each with its own memory space, and takes 17 min 39 s. The fastest hybrid case places 10 processes on each node with 4 threads per process and takes 15 min 06 s. The number of independent memory spaces on a node therefore falls from 40 to 10, and the threads within a process share that process's memory. The supplement already states that each pure-MPI process holds an independent memory space, whereas the hybrid mode runs threads in shared memory and thereby eases contention in intra-node memory access. Lines 567-569 will state the reduction from 40 independent memory spaces per node to 10 as the basis for the lower memory pressure.

Comment 19:

I feel all the figures need another pass for readability. Several maps lack coordinates or scale bars, and some colorbars and line styles are not explained. Please explicitly mark the model ignition point or line on a georeferenced map. The time axes in Fig. 7 should identify UTC, and the "Area" values in Fig. 5c should be labelled as overlap areas.

Response to comment 19:

A further pass over the figures is specific, and the difficulty in reading them is regretted. The geographic figures, including the burned extent, the nested domains and the fire-area maps, will be given latitude and longitude coordinates and scale bars. The quantity and units of each colorbar, and the experiment represented by each line style, will be stated in the legend or the caption. The model uses an ignition line, and that line will be marked on a georeferenced fire-area map at the location given in the text, 119.09°E, 49.20°N. The time axis of Fig. 7 will be labelled UTC. The Area values in Fig. 5c are the overlap areas used to compute recall, recall being that overlap divided by the surveyed 1299 km². In Fig. 5c, Area will be relabelled as overlap area, in keeping with the definition of recall in the caption.

Comment 20:

The text, tables and references need a thorough proofread. For example, "utilizationefficiency" at L482, the Fig. 2 caption is repeated in the text at L153.

Response to comment 20:

The text, the tables and the references will be proofread throughout. At line 482, utilizationefficiency will be changed to utilization efficiency. At line 153, the caption

of Fig. 2 was carried into the sentence. That sentence will keep only “A three-level nested grid system was used for the numerical simulations (Fig. 2)”, and the repeated caption will be removed.